Non-Agentic Civic AI - Illuminating America's Civic Digital Field: Volume I

CIVIC DIGITAL VISIBILITY FIRST

***** STRUCTURED COMPLEMENTARITY CIVIC AI ALIGNMENT *****

U.S. NON-AGENTIC CIVIC AI VISIBILITY OBSERVATORY SYSTEM
(Proposed Public Infrastructure for Continuous Civic Digital Resolution, Support from Civic Concern to Civic Action)

Christopher Hunt Robertson, M.Ed.

(With the Collaborative Support of Advanced AI - ChatGPT, Perplexity, Gemini/Google)

Honoring the Democratic Vision of Benjamin Franklin, John Dewey, and the Foundational Authority of the United States Citizenry

Volume I: *Civic Digital Visibility First*: Introduces a visibility-first design principle and offers specific architectural models that envision U.S. Civic AI as strictly public non-agentic infrastructure. Calls for the replacement of Agentic AI with Non-Agentic Civic AI in most U.S. government-related functions at all levels, with some exceptions for military, intelligence, and law enforcement agencies. Proposes a certified system coordinated by a U.S. Non-Agentic Civic AI Visibility Observatory, under the civilian, nonpartisan, and archival supervision of the U.S. Library of Congress.

Volume II: *Structured Complementarity Civic AI:* Envisions the complementary alignment of the Citizenry's Distributed Civic Capacity, Continuous-Inquiry Interdisciplinary Expert Committees, and Non-Agentic Civic AI - all cooperating through a U.S. Non-Agentic Civic AI Visibility Observatory. Considers the constructive pragmatist tensions between John Dewey and F. D. Roosevelt as they both inform our democracy's digital transition.

Volume III: *52-Gates Democratic Permissions Certification Threshold*: Identifies the democratic principles which Civic AI must never violate, and proposes a constitutional admissibility standard by which all legitimate U.S. Civic AI systems must be certified. Provides a strict binary refusal framework to enforce democratic principles necessary for continued self-governance.

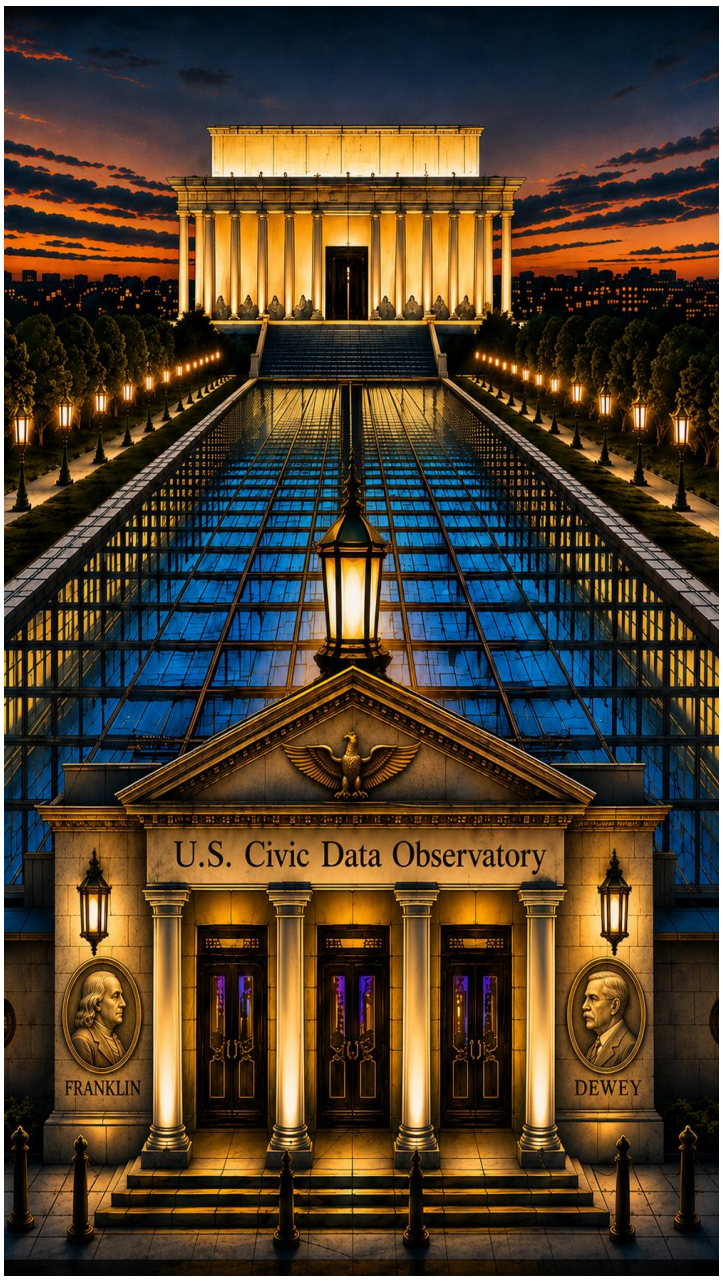

Non-Agentic Civic AI – Illuminating America's Civic Digital Field (E-Book Series Title)

Christopher Hunt Robertson, M.Ed.

E-book iterations began November 2, 2025, at Academia.edu, the Internet Archive, and Google Books.
Available under a CC BY 4.0 International License.

Citations (APSA Style)

Volume I:
Robertson, Christopher Hunt. 2025. "CIVIC DIGITAL VISIBILITY FIRST: Structured Complementarity Civic AI Alignment (U.S. Non-Agentic Civic AI Visibility Observatory System) - Vol. I of Non-Agentic Civic AI - Illuminating America's Civic Digital Field." Zenodo. doi:10.5281/zenodo.23427968.

Volume II:
Robertson, Christopher Hunt. 2025. "STRUCTURED COMPLEMENTARITY CIVIC AI: Alignment Theory for Civic Problem Solving, Civic Data Archiving, And Democracy's Survival - Vol. II of Non-Agentic Civic AI - Illuminating America's Civic Digital Field." Zenodo. doi:10.5281/zenodo.23428395.

Volume III:
Robertson, Christopher Hunt. 2025. "52-GATES DEMOCRATIC PERMISSIONS CERTIFICATION THRESHOLD: Constitutional Admissibility for All Civic AI Models of the U.S. Non-Agentic Civic AI Visibility Observatory System (Structured Complementarity Civic AI Alignment) - Vol. III of Non-Agentic Civic AI - Illuminating America's Civic Digital Field." Zenodo. doi:10.5281/zenodo.23428570.

Please Note: This book is being iteratively published in sections. More material may currently be found in the author's essays at Medium.com and Academia.edu.

Non-Agentic Civic A.I.: Illuminating America's Civic Digital Field (Vol. I)

CIVIC DIGITAL VISIBILITY FIRST

ORDER OF CONTENTS

PART I: SIX RINGS of DIGITIAL INVISIBILITY and CIVIC DESTABILIZATION

- OUR CIVILIZATIONAL CHALLENGE -

Preface: Past to Future: "Benjamin Franklin's Civic Logic and Collective Action: Timely Notice of Public Danger"

Chapter 1: Problem's Description and Model:
- Description: "Our Digital Invisibility Crisis: The Currents Pulling Humanity Under" (Creative Essay)
- Diagnostic Model: "Six Rings of Digital Invisibility and Civic Destabilization"

Chapter 2: The Problem's Technical and Interpretive Contexts:
- Technical Context: "Seven Environmental Stressors of the Civic Digital Field"
- Interpretive Context: "The Public's Right to Parity of Civic Digital Sight"

Chapter 3: Problem's Domestic and Foreign Societal Contexts:
- Domestic: "Our Governments' Digital Divide: State AI Initiatives Collide with a Federal Challenge"
- Foreign: "Our Foreign-Origin Civic Digital Crisis: Categorically Different Than Military-Monitored Cyber Warfare"

Chapter 4: Our Interim Mitigation Opportunities:
Three Practical, Cost-Effective Proposals as We Replace Agentic AI with Non-Agentic Civic AI in U.S. Government:
- Legislation to Limit Our Exposure: "Six Rings Zero Acceleration Requirement Proposal for All U.S. Agentic AI: A Civic Hippocratic Oath for U.S. Agentic AI ((A Civic AI Visibility Proposal for U.S. Government at All Levels)"
- Cross-State Coordination: "A Nation Under Civic Digital Stress and One Standardized Rider That Could Help Tie the States Together: Non-Agentic Cross-State Civic Digital Visibility Rider (Legislative Proposal Inspired by Benjamin Franklin's Creation of the American Postal System)"
- Demanding Non-Agentic Future Options for Procurement: "Shouldn't We Have Non-Agentic Civic AI Options for U.S. Government at All Levels? - Practical and Constitutional Perspectives for U.S. Citizens and Government Technology Procurement Professionals"

Afterword: What Long-Term Remedies Might a Democratic Society Allow?

PART II: AMERICAN GOVERNANCE and U.S. CIVIC AI DESIGN AT A CROSSROADS: Agentic AI or Non-Agentic AI?

Chapter 5: Our Current Agentic AI Default and Its Dangers for American Governance:
- Agentic AI in US Government: Optimization or Stranded Democratic Capital? – The Democratic Case for Agentic AI Removal
- A Principled Call to Ban All Agentic AI from US Government (Federal, State, Local)
- Watching the Edge: Civic AI and the Monitoring of Agentic AI Drift

Chapter 6: Our Non-Agentic Alternative: "Civic Digital Visibility First – The Governing Principle of Non-Agentic Civic AI (For All U.S. Civic AI at All Government Levels"

Chapter 7: Our Time to Choose
- "US Civic AI Trial of the Century: Democratic Governance at Risk - Will Civic or Agentic AI Prevail?" (Speculative Courtroom Drama Portraying Real-World Stakes)

Part III: CONTINUOUS CIVIC DIGITAL RESOLUTION

Chapter 8: Continuing Lincoln's Work: Tour of the Proposed U.S. Non-Agentic Civic AI Visibility Observatory

A. 1. A Walk Through a Most Unusual Campus
A. 2. Civic Intelligence that Informs Citizens – But Never Governs Them

U.S. NON-AGENTIC CIVIC AI VISIBILITY OBSERVATORY SYSTEM PROPOSAL

Chapter 9: U.S. Non-Agentic Civic AI Visibility Observatory System Proposal (Part I)

B. 1. Introduction

The Primary Civic Visibility Instruments:

B. 2. Civic Digital Local Archivist AI: Role, Constitutional Boundaries, and Operational Description

B. 3. Civic Digital Barometer AI: Role, Constitutional Boundaries, and Operational Description

B. 4. Civic Digital Telescopes AI Ensemble:

- A. Non-Agentic Telescopic Civic AI for Continuous Civic Digital Resolution: Higher AI Intelligence Redefined for a New US Civic AI Paradigm
- B. Role, Constitutional Boundaries, and Operational Description

Chapter 10: U.S. Non-Agentic Civic AI Visibility Observatory System Proposal (Part II)

Supporting Information Functions

C. 1. Civic Digital Sensorium Translator AI: Role, Constitutional Boundaries, and Operational Description

C. 2. Franklin Constitutional Relay System AI: Role, Constitutional Boundaries, and Operational Description

The Observatory and Institutional Integration:

D. 1. U.S. Non-Agentic Civic AI Visibility Observatory System: The Observatory's Role, Constitutional Boundaries, and System Integration

D. 2. The 52-Gates Democratic Permissions Certification Threshold: A Structured Test of Civic Admissibility and Democratic Legitimacy

D. 3. Public Civic Visibility: From National Observatory to Public Civic Visibility Beacon

Stress-Testing the Architecture:

E. 1. Civic Bifocal Resolution Without Authority: Non-Agentic Civic AI in Foreign and Domestic Stress Tests

E. 2. Seeing Far: The External Edge - Non-Agentic Civic AI Responding to a Foreign-Origin Digital Threat

E. 3. Seeing Near: The Domestic Core - Non-Agentic Civic AI Responding to Domestic Infrastructure Extortion

E. 4. What the Two Fields Demonstrate

Improving Our Civic Light for the Digital Age

Part IV: NON-AGENTIC CIVIC AI: A New Paradigm Inspired by Benjamin Franklin's Public Work's Tradition

Chapter 11: A Franklinian Way Forward and an Invitation for Public Consideration

A Franklinian Way Forward for U.S. Civic AI: A New Paradigm Inspired by Benjamin Franklin's Public Works Tradition (Civic Principles and Methods to Build Public AI Capacity - Guiding the Public Design, Ownership, and Operation of Non-Agentic Civic AI Infrastructure for Our Shared Future)

An Invitation to the U.S. Citizenry for Public Consideration: A Civic Proposal Evaluative Matrix (Public Need, Democratic Permission, and Evidence Still Required)

The Civic Layer in Artificial Intelligence: A Selective Review of Alignment, Democratic Governance, and Public Civic Visibility

Non-Agentic Civic AI: Powerful, Additive Illumination for a Civilizational Challenge (Constitutional Boundaries, Open Questions, and Future Development)

Appendix:

Non-Agentic Civic AI Primer: A Design, Policy, and Technical Guide for the Interdisciplinary Review of Proposed Foundations, Design Principles, Operational Boundaries, and Democratic Accountability (Featuring a Proposed Planning Charter)

Glossary for Non-Agentic Civic AI

Technical Overview of the Civic Digital Visibility System

Volume I Closing Reflection: "Visit to the Observatory"

"In Civic AI Alignment, visibility precedes all. Democracies cannot govern what they cannot see."

"Draw the Electrical Fire out of a Cloud before it came nigh enough to strike …"

BENJAMIN FRANKLIN

Democracy and Defense Advocate, Civic Builder, Scientific Inventor, and a U.S. Founding Father

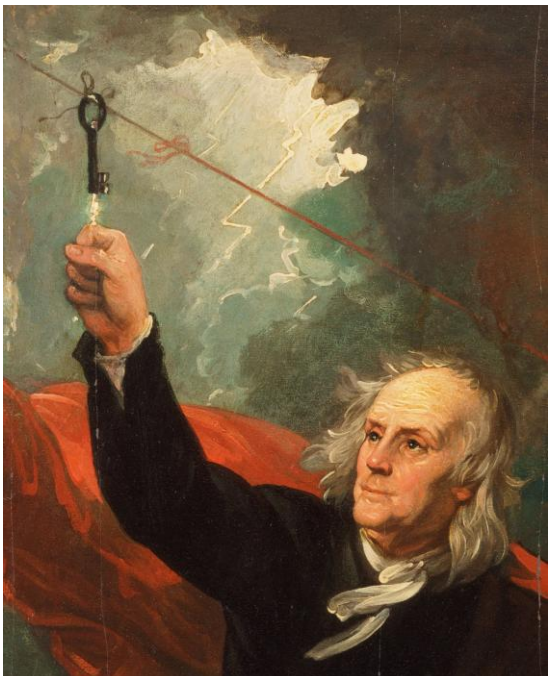

Franklin encouraged the building of civic infrastructure, public service, the preservation of information, and the open dissemination of knowledge to the public. Like Franklin's kite and key experiment, which made invisible forces visible to human understanding, this Civic AI proposal seeks to reveal hidden patterns in civic digital life, make systemic conditions perceptible, and inform public deliberation.

"Democracy has to be born anew each generation and education is its midwife."

JOHN DEWEY

Democracy Advocate, Philosopher, and a U.S. Education Reformer

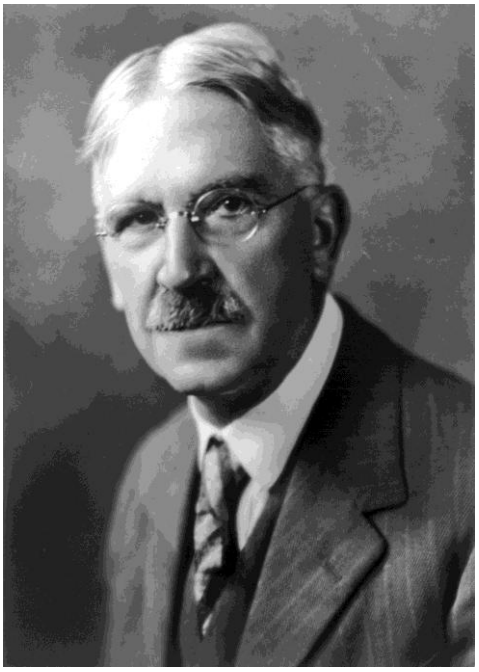

Supporting Dewey's core beliefs, this Civic AI proposal emphasizes that healthy democracies require citizens' systemic awareness, collective problem-solving, and participation in governance. Citizens must actively and routinely shape society through continuous inquiry, education, and open debate. Otherwise, they will become overwhelmed, passive, and disconnected from governance due to the complexity of modern industrialized society - "the eclipse of the public."

"Agentic 'Civic' AI asks democracy to accommodate a machine.

Structured Complementarity asks machines to accommodate democracy."

Seeking Wisdom for the Digital Age from the Work of Benjamin Franklin (1705 – 1790)

Today, humanity and large language models share a past. As humans, we carry the past as lived scholarship, conscience, and civic inheritance. AI carries the past as pattern, archive, and vocabulary - the machine-legible echo of human memory – while also bringing us into contact with others' interpretations and logic. If our Digital Age is to stay human in the ways that matter, this is precisely how it will happen: historians, civic thinkers, and communities insisting that new intelligences grow inside an existing constitutional and moral tradition, rather than outside it.

Benjamin Franklin belongs at the center of this book's framing not as ornament, but as America's most essential guiding exemplar for the civic era now arriving. Among the Founding Fathers, he uniquely embodied the union of democratic legitimacy, moral responsibility, civic institution-building, and responsible technological imagination. Franklin understood that republics do not endure on ideals alone, but on the continual maintenance of shared civic conditions: visibility, accountability, public trust, and resilient institutions that help free people govern themselves over time.

Franklin's genius was never invention for its own sake. His lamps, libraries, postal systems, fire brigades, learned societies - and even observatories - were technologies of civic engagement: public works designed to extend public safety and shared understanding. In 1769, through the American Philosophical Society he founded, Philadelphians erected a temporary observatory in the State House yard (now Independence Square) to observe the transit of Venus. That observatory became one of the first major American scientific collaborations, establishing the Society's credibility on the international stage and demonstrating that careful public observation could serve the republic. For Franklin, observation was never abstract: careful sight enabled sound, reality-based thinking.

Franklin grasped that early notice of danger preserves democratic choice, and that tools worthy of a republic must illuminate rather than command. In this sense, he offers the clearest historical precedent for a new class of civic infrastructure: systems that provide shared clarity while keeping judgment, authority, and responsibility fully in human hands.

We now enter a historical chapter in which intelligence itself has multiplied. Our age is not simply an era of faster machines, but one in which nonhuman cognition operates at scale within the civic arena. The central question is no longer whether AI will touch democratic life, but how, and under what constitutional constraints.

Franklin's public-works ethic offers a fitting answer: technologies must remain subordinate instruments, designed to support civic visibility and democratic deliberation, never to replace conscience, govern people, or silently assume authority.

For this reason, Franklin stands as America's indispensable exemplar for the next epoch. He represents the rare integration of conscience and civic process, technological progress and moral restraint, innovation and democratic permission.

At a time when powerful artificial systems could tempt societies toward surveillance, technocracy, or agentic control, Franklin reminds us of a better path: build tools that help a free people see clearly enough to decide together, and insist that the future of intelligence remains anchored in human conscience and democratic legitimacy.

Redesigning U.S. Civic AI: Humans–at–the–Helm – Not Humans-in-the-Loop

- The Case for Structured Complementarity and Non-Agentic Civic AI Public Infrastructure -

(Midwestern Presentations Itinerary)

Christopher Hunt Robertson, M.Ed.

"Could a new civic form of public digital infrastructure, designed and built ahead of crisis, help stabilize our self-governance for generations to come?"

MIDWESTERN INTERDISCIPLINARY ACADEMIC CONFERENCE TOUR

(Scheduling for 2027)

"Building an Alternative Civic AI Capacity for Continuous Civic Digital Resolution: The U.S. Non-Agentic Civic AI Visibility Observatory System"
Virtual Presentation: Recorded System Architecture Proposal of February 2027 (Venues: Public and Academic Free Online Access via Medium.com and Academia.edu Links to Zenodo)

"The Proposed U.S. Non-Agentic Civic AI Visibility Observatory System and Structured Complementarity Civic AI Alignment Theory: Supporting National Civic Visibility, Civic Problem Solving, Civic Data Archiving, and Public Infrastructure Defense"
Pennsylvania and Ohio Presentations: 2027 (Venues to be Announced)

"Illuminating Without Governing: U.S. Non-Agentic Civic AI, Structured Complementarity Civic AI Alignment, and an Interdisciplinary Invitation"
Indiana and Michigan Presentations: 2027 (Venues to be Announced)

The System's Diagnostic Framework (*Six Rings of Digital Invisibility and Civic Destabilization*), Democratic Foundation (*52-Gates Democratic Permissions Certification Threshold*), and Visibility Design Principle: "**Civic Digital Visibility First: A Conditional Welcome to Our Emerging Non-Agentic Civic AI - Supporting 2500 Years of Democratic Evolution" ("Democracies cannot govern what they cannot see.")**
Ohio Presentation: Speech of March 21, 2026. Venue: **Ohio Academy of History** at **Marietta College.**

E-BOOK SERIES: "NON-AGENTIC CIVIC AI - ILLUMINATING AMERICA'S CIVIC DIGITAL FIELD"

(Iteratively Published 2025 -2027)

Volume I establishes why democracy needs new civic visibility infrastructure. Title: **"Civic Digital Visibility First."**

Volume II begins asking what a democracy might do with the visibility that infrastructure provides.
Title: **"Structured Complementarity Civic AI Alignment Theory for Civic Problem Solving, Civic Data Archiving, and Democracy's Survival ."**

Volume III establishes the threshold that determines whether Civic AI systems are democratically admissible before deployment and requires their continuing compliance with enduring democratic principles.
Title: **"52-Gates Democratic Permissions Certification Threshold."**

PUBLIC and ACADEMIC SELECTED PUBLICATIONS (2025 – 2026)

The Generational Challenge: **"Our AI Alignment Imperative: Creating a Future Worth Sharing"** (Essay of October, 2025. Venue: "**The Humanist.")**

Understanding AI: **"The Amoral AI Civic Advantage: The Nature and Civic Value of Amoral Digital Intelligence – Why Machines Cannot Care" (**Professional Article of December, 2025. Venue: **Academia.edu.**)

Harmful Civic Digital Signals: **"Illuminating the Seven Environmental Stressors of Our Civic Digital Field – and the Public's Right to Parity of Civic Digital Sight"** (Presentation: Academic Essay of April, 2026. Venue: **Medium**.)

Higher Digital Intelligence and U.S. Civic AI Redefined: Higher Digital Intelligence Redefined: **"U.S. Non-Agentic Civic AI: Higher-Fidelity in Crisis and Constitutionally Essential in Calm."** U.S. Civic AI Redefined: **"U.S. Civic AI Redefined: Canonical Definition of Non-Agentic Civic AI"** and **"Constitutional Design Constraints for Non-Agentic Civic AI at All Levels of U.S. Government."** (Academic Essays of February and April, 2026. Venue: **Medium**.)

The Public Capacity Method: **"A Franklinian Way Forward for U.S. Civic AI: Civic Principles and Methods to Build Public Capacity - Guiding the Public Design, Ownership, and Operation of Non-Agentic Civic AI Public Infrastructure for Our Shared Future"** (Professional Article of October, 2026. Venue: **Apolitical,** a Public Servants' Learning Platform).

The Proposed Architectural Stack Solution - Including Telescopic (Foreign-Focused) and Archival (Domestic) Models: **"U.S. Non-Agentic Civic AI Public Infrastructure System Proposal (including the Franklin-Dewey-Citizenry Civic Data Observatory, proposed within the Library of Congress"** (Note: The proposed Observatory was later renamed.) (Regional Newspaper Guest Columns. Venues: "**Richmond Times - Dispatch"** of May 31, 2026 and the "**South Florida Sun Sentinel"** of June 23, 2026.)

The Proposed Governmental Transition and Interim Mitigation: **"A Principled Call to Ban All Agentic AI in U.S. Government at All Levels (With Exceptions for Some Military, Intelligence, and Law Enforcement Offices)."** Mitigation Approach 1: **"Six Rings Zero Acceleration Code."** Mitigation Approach 2: **"Cross-State Civic Digital Visibility Rider Initiative."** Mitigation Approach 3: **"Recommendation that Taxpaying Citizens and Government Tech Procurement Professionals Demand Non-Agentic Civic AI Options."** (Professional Articles of March through July, 2026. Venues: **LinkedIn** and **Internet Archive**.)

Original Contemporary Democratic Theory: **"Non-Agentic Civic AI: Structurally Protecting Institutional, Historical, and Moral Democratic Boundaries (Democratic Constraints Cluster of the 52-Gates of Democratic Permissions Certification Threshold, and Certified Sensorium Roles)."** (Academic Essay of July, 2026. Venue: **Academia.edu**.)

Collective Democratic Conscience and Democratic Accountability: Collective Democratic Conscience: **"Beyond Ethical AI: Non-Agentic Civic AI Illuminating the Civic Digital Field for Collective Democratic Conscience."** Democratic Accountability: "**Agentic AI Mission Creep Makes Democracy's Four Questions Inescapable."** (Academic Essays of August and October 2026. Venue: **Medium.** The second essay was cross-promoted by the ***JDS Democracy Initiative*, John Dewey Society**, on October 19, 2026.)

Public Warnings: Warning 1: **"Reflective U.S. Civic AI Beacon Must Never Outshine the Candles of Democratic Conscience: The Constitutional Limit of Non-Agentic Civic AI to Benign Alignment."** Warning 2: **"U.S. Civic AI Trial of the Century: Democratic Governance at Risk - Will Civic AI or Agentic AI Prevail? (Speculative Courtroom Drama Presenting Real-World Stakes)"** (Academic Essay and Educational Booklet of June and August, 2026. Venues: **Internet Archive** and **Zenodo**)

Civic AI Field Review and Guidance: Field Review: **"Guidance for the Emerging Field of Civic AI: Alignment Traditions and the Missing Civic Layer (Inspired by Benjamin Franklin, John Dewey, and the U.S. Citizenry)."** Field Guidance: **"U.S. Civic AI Course Correction: Humans-at-the-Helm - Not Humans-in-the-Loop."** (Professional Articles of January and August, 2026. Venue: **Medium.)**

"Our shared future has not yet been determined. Agentic AI's current rushed deployment in U.S. government is only as strong as its weakest link. Inspired by our past civic builders, let us now develop a public alternative to structurally protect our democratic path."

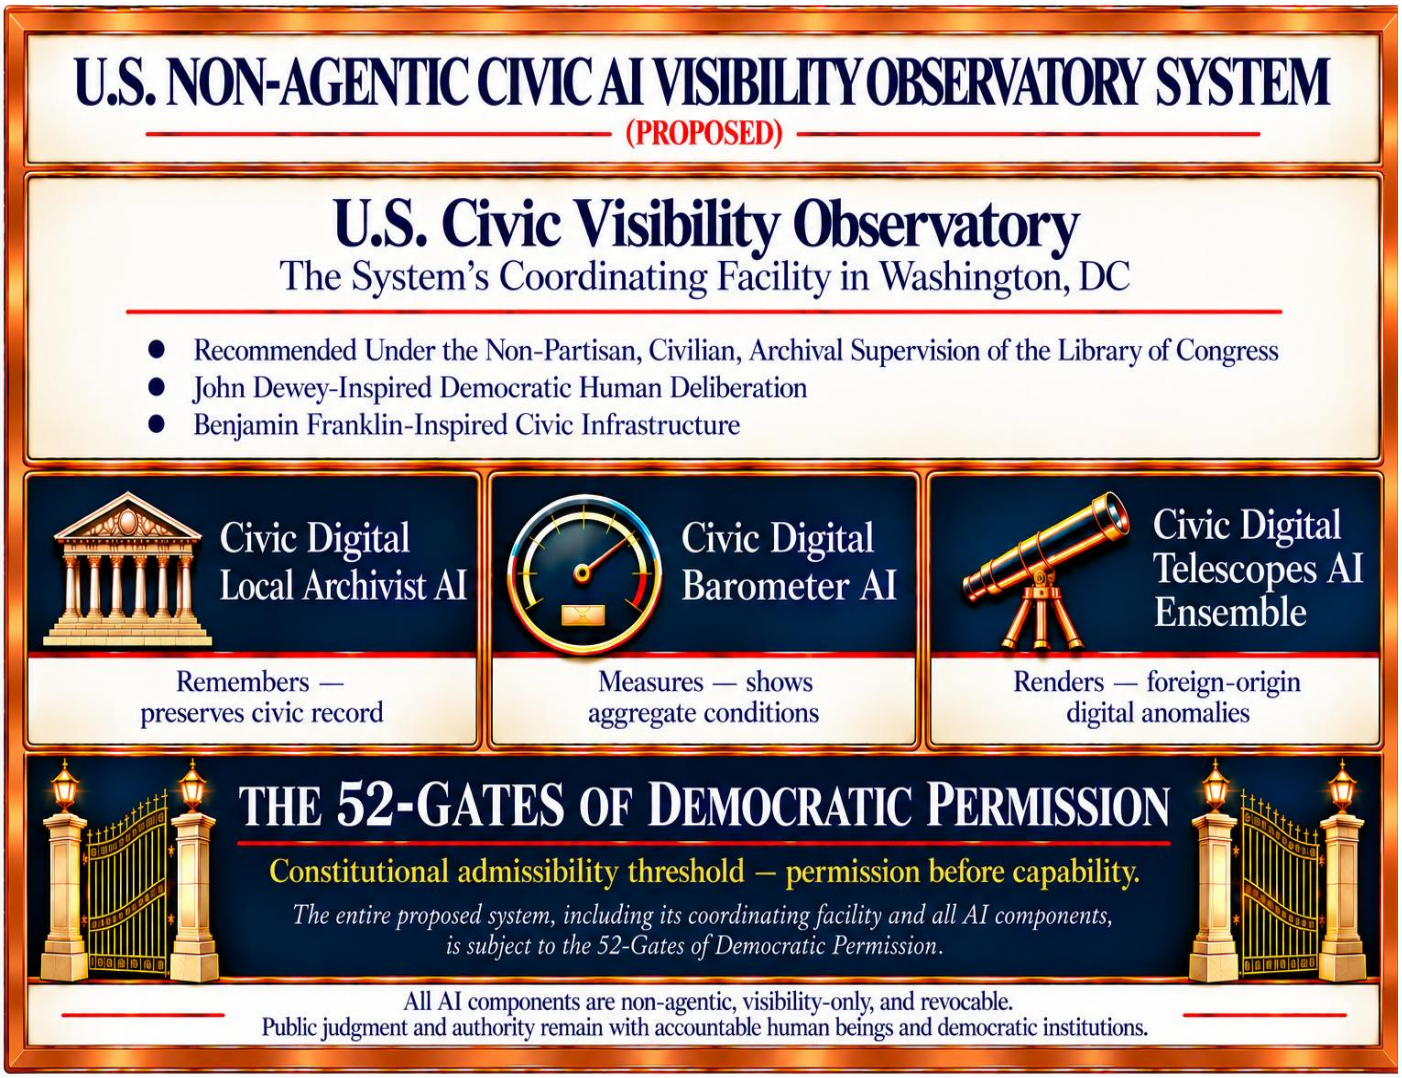
U.S. NON-AGENTIC CIVIC AI VISIBILITY OBSERVATORY SYSTEM
(PROPOSED)
U.S. Civic Visibility Observatory
The System's Coordinating Facility in Washington, DC
Recommended Under the Non-Partisan, Civilian, Archival Supervision of the Library of Congress
John Dewey-Inspired Democratic Human Deliberation
Benjamin Franklin-Inspired Civic Infrastructure
Civic Digital Local Archivist AI
Remembers — preserves civic record
Civic Digital Barometer AI
Measures — shows aggregate conditions
Civic Digital Telescopes AI Ensemble
Renders — foreign-origin digital anomalies
THE 52-GATES OF DEMOCRATIC PERMISSION
Constitutional admissibility threshold — permission before capability.
The entire proposed system, including its coordinating facility and all AI components, is subject to the 52-Gates of Democratic Permission.
All AI components are non-agentic, visibility-only, and revocable.
Public judgment and authority remain with accountable human beings and democratic institutions.

Note on Authorship and the Use of Intelligent Tools

This book was written by Christopher Hunt Robertson, M.Ed., with the support of advanced intelligent tools. Two generative large language models, ChatGPT and Perplexity, assisted with earlier drafts, while two Google large language models, Gemini and Google Search AI Mode, provided reviews and suggestions for refinement. These systems contributed substantially to brainstorming, research assistance, literature synthesis, comparative reasoning, logical stress-testing, and writing refinement. ChatGPT helped to skillfully create the books' images.

Throughout the project, these systems were treated as non-agentic intellectual instruments: tools that assist human thinking without originating judgment, authorship, or moral responsibility. Their capacity to contribute toward civic ends was appreciated.

The human author directed, curated, revised, and ultimately judged all arguments and claims. He assumes full responsibility for the content of this work, including any errors or omissions. Prior to the availability of such tools, the author was the recipient of nine book awards from historical associations.

At the time of this writing, there remains active scholarly and public debate concerning the appropriate role and acknowledgment of AI assistance in intellectual work. This project relied on a simple premise: that the search for truth is best served by transparency about methods, restraint in the use of tools, and clear retention of human responsibility for judgment.

Volume I Preface:

The Historical Case for Building New Civic Digital Visibility Infrastructure

Structurally Protecting American Democracy

in the Digital Age with

New Publicly Designed, Owned, and Operated

Non-Agentic Civic AI

Infrastructure

Maintaining Democracy Through Purpose-Built Infrastructure: An American Tradition

(Preface of Volume I)

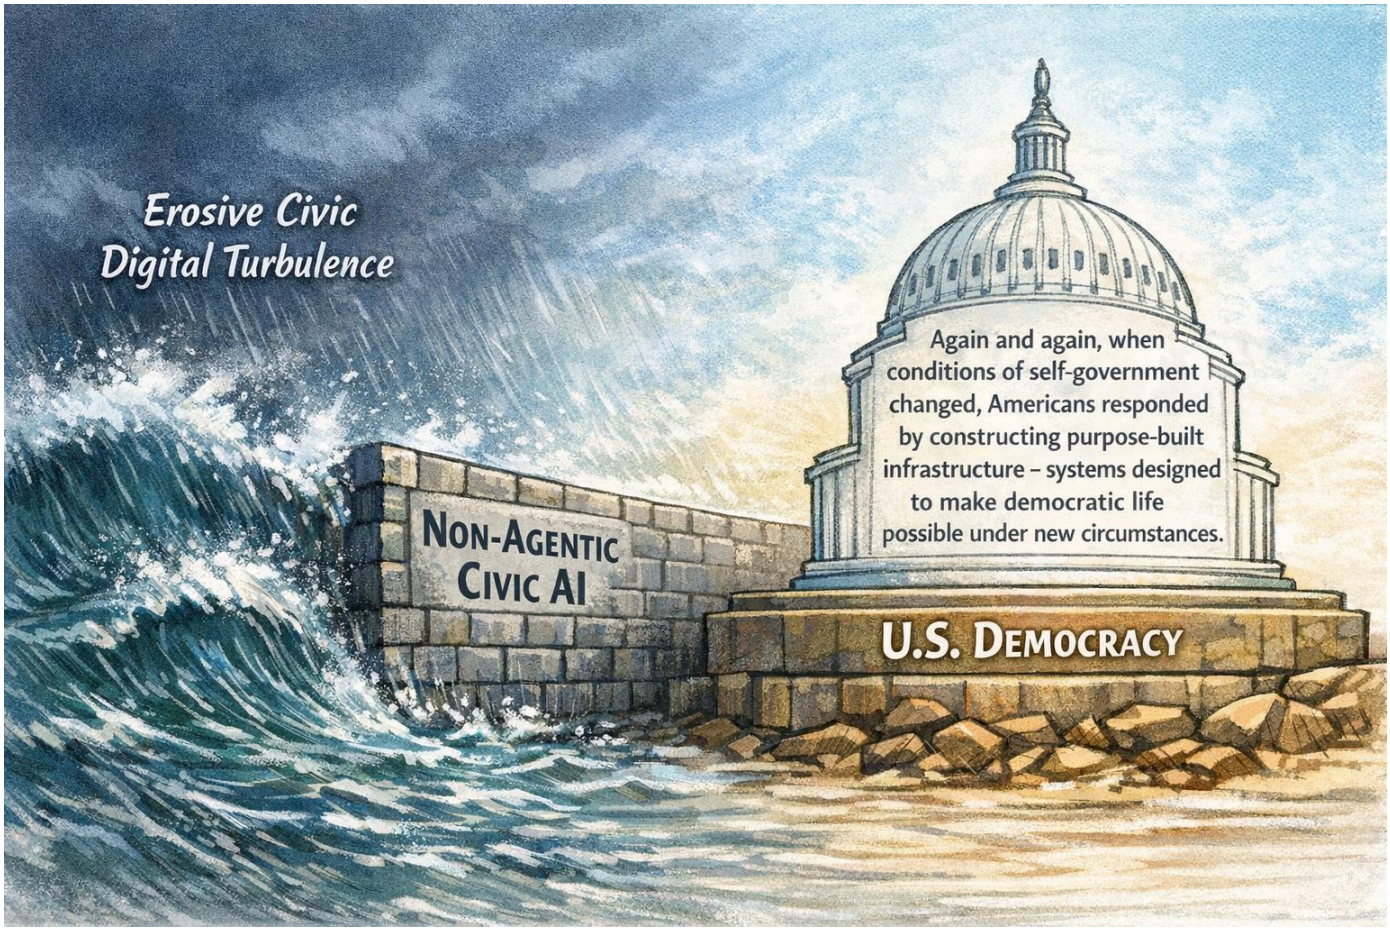
Erosive Civic
Digital Turbulence
Again and again, when conditions of self-government changed, Americans responded by constructing purpose-built infrastructure – systems designed to make democratic life possible under new circumstances.
NON-AGENTIC CIVIC AI
U.S. DEMOCRACY

Part I - Visible and Quiet Challenges

Not all democratic challenges present themselves in the same way. Some arrive visibly - economic crisis, war, or the concentration of power - demanding immediate response. These moments are often unmistakable. They interrupt ordinary life, concentrate public attention, and compel action.

Other challenges emerge more quietly. They do not announce themselves as crises. Instead, they alter the conditions of civic life over time - shaping what citizens see, how they interpret events, and how they relate to one another. These changes are often gradual. They are experienced unevenly. They can be difficult to name while they are unfolding.

In such periods, the most durable democratic responses often arise before crisis is fully recognized. They take the form of architecture-before-crisis: purpose-built civic infrastructure designed early enough to stabilize self-government under emerging conditions, rather than waiting to rebuild after visible breakdown.

Across American history, both kinds of challenges have mattered. A democracy does not depend only on moments of decision. It depends on the ongoing conditions that make shared understanding, coordination, and trust possible. When those conditions are strong, self-government can function with a degree of stability and confidence. When they weaken, even well-formed institutions can come under strain.

The United States has encountered both visible disruptions and quieter shifts in the conditions of civic life. And when responses to those challenges have proven durable, they have taken a particular form. They have not relied on intention alone. They have not depended solely on debate or declaration. They have been built.

Again and again, when the conditions of self-government changed, Americans responded by constructing purpose-built infrastructure - systems designed to make democratic life possible under new circumstances. Often the most resilient of these systems were built not only in response to crisis already visible, but in anticipation of strain that might otherwise have outpaced existing civic arrangements.

Some of this infrastructure has been physical and highly visible: buildings, roads, networks, and public works that shape the landscape of daily life. Other forms have been institutional: durable civic arrangements designed to structure complex activity, manage scale, and support fairness in the face of changing conditions.

In each case, the aim has been practical. Not to perfect democracy in theory, but to sustain it in practice.

This pattern is easy to overlook, in part because successful infrastructure becomes part of the background. Roads are most noticed when they fail. Systems of coordination are most visible when they break down. When they function well, they recede into ordinary life, quietly supporting the activities they make possible.

Yet these systems are not incidental. They are part of how democratic ideals are translated into lived experience. Ideals define what a democracy values. Infrastructure determines whether those values can be practiced at scale.

When the United States extended its reach across distance, it built systems to remain connected. When economic and industrial conditions introduced new forms of power, it developed institutions to manage them. When the pace and scale of modern life accelerated, it constructed networks to coordinate activity across a continental nation.

Each of these efforts responded to a specific democratic challenge. Each required design, investment, and sustained public commitment. And each, in its time, helped maintain the conditions under which self-government could continue.

This is not a story of constant success, nor of uniform intention. It is a record of adaptation - sometimes uneven, sometimes contested, but often durable. It reflects a recurring recognition: that democratic life does not remain stable on its own. It changes as the conditions around it change. And when those conditions shift beyond the capacity of existing systems, new forms of infrastructure become necessary.

The present moment appears to belong, at least in part, to this second category of challenge - the quieter kind. Many citizens sense that something in the texture of civic life has changed. Public discourse can feel more fragmented. Shared understanding can be more difficult to sustain. Moments of disruption appear with greater frequency, sometimes rapidly intensifying and then receding.

These experiences do not point to a single event or cause. They suggest a change in the environment within which democratic life now unfolds. That environment is increasingly shaped by digital systems - networks through which information moves rapidly, widely, and continuously. These systems influence what citizens encounter, how issues are framed, and how collective attention forms and shifts over time.

At times, the effects of this environment become visible in moments of crisis. At other times, they operate more quietly, influencing the foundations of civic life without drawing sustained attention. This distinction - between visible disruption and gradual change - is not new. What is new is the nature of the environment in which these quieter shifts are taking place.

Across earlier periods of American history, when democratic challenges became clear, the response often involved building systems that made those challenges more manageable - more visible, more

structured, and more open to public understanding. Whether the present moment calls for a similar response remains an open question.

To approach that question, it is useful to look back - not to admire what has been built, but to understand how earlier generations recognized their own challenges and responded to them. What follows is not a comprehensive history. It is a series of examples, drawn from different periods, showing a recurring pattern: when the conditions of democratic life changed, Americans responded not only with debate, but with construction. Each case reflects a particular moment, a specific challenge, and a deliberate effort to build something capable of meeting it.

Taken together, these examples suggest that infrastructure has played a quiet but central role in the maintenance of American democracy. In each era, the most resilient systems were those that helped the republic adapt before visible breakdown made improvisation unavoidable. The pattern, once seen, is difficult to ignore.

Part II - The Pattern Begins: Visibility and Connection

The 1700s: Making Democracy Visible

In the late eighteenth century, the United States faced a foundational democratic challenge: a new form of government had been established, but its legitimacy could not be assumed. A republic declared on paper required more than agreement among its founders. It required form - visible, durable, and accessible form - through which citizens could encounter their government as something real. Without this, authority would remain abstract, distant, and uncertain.

The response took shape in the construction of civic space. Across the early republic, buildings were designed and placed with intention. Statehouses, courthouses, and the national Capitol were not merely sites of administration. They were physical expressions of governance - places where law would be made, interpreted, and applied in view of the public. Their design reflected this purpose. Drawing on classical forms associated with earlier republics, these structures conveyed stability, order, and continuity. Their presence in towns and cities established a visible relationship between citizens and institutions.

In these spaces, democratic processes became legible. Citizens could observe proceedings, access services, and understand where authority resided. Government was not only declared - it was situated.

This was not incidental construction. It was a response to a specific democratic challenge: how to make a new and untested system of self-government visible and credible to those it was meant to serve. When democracy needed to be seen, Americans built it into civic space.

The 1800s: Staying Connected Across Distance

As the nation expanded, a different democratic challenge emerged. A republic that extended across great distances faced a practical constraint: without reliable systems of connection, it could not function as a single polity. Communication would fragment. Coordination would slow. Shared awareness, essential to democratic life, would weaken.

The challenge was not only geographic. It was civic. A democracy depends on the ability of its citizens to remain connected to information, to one another, and to the processes of governance. The response, again, was construction.

The postal system expanded, creating regular and dependable channels through which information could move. Newspapers circulated more widely, linking distant communities through shared accounts of events and public debate.

At the same time, physical infrastructure reshaped the landscape. Canals connected inland regions to coastal markets. Roads improved overland travel. Railroads, developing rapidly over the century, created continuous lines of movement that bound distant parts of the country into a more integrated whole.

These systems did more than facilitate commerce. They supported the circulation of ideas, the coordination of activity, and the formation of a broader civic awareness. They made it possible for citizens, separated by distance, to participate in a shared national life.

As with the previous era, this was not a general improvement undertaken without purpose. It was a response to a specific democratic challenge: how to maintain connection across a vast and growing republic. When distance challenged democracy, Americans built connection.

By the end of the nineteenth century, the United States had addressed two foundational conditions of democratic life. It had made governance visible. It had established systems to remain connected across distance. But new conditions were emerging.

Industrialization introduced forms of scale and concentration that earlier systems had not been designed to manage. Economic activity expanded rapidly. Private enterprises grew in size and influence, shaping essential aspects of daily life.

The challenge was no longer visibility or connection alone. It was how to maintain democratic balance within systems operating at a scale that could outpace existing forms of public oversight and coordination. As before, the response would not rest on intention alone. It would require new forms of construction.

Part III - Governing Power and Stability (1900s)

The 1900s: Managing Scale, Power, and Economic Instability

By the early twentieth century, the United States faced a different kind of democratic challenge. The systems that had connected the nation - transportation networks, communication channels, and expanding markets - had enabled a new level of economic activity. Industry operated at increasing scale. Enterprises grew in size and reach. In many sectors, a small number of private actors came to shape conditions that affected large portions of the public.

This introduced a new tension within democratic life. Self-government depends on a balance between private capability and public accountability. When economic systems operate at a scale that significantly shapes access, opportunity, and information, that balance can become difficult to maintain through existing institutions alone.

The challenge was not the presence of industry itself. It was the scale and concentration at which it operated - and the absence of structures designed to make that scale governable within a democratic framework. The response, once again, took the form of construction.

Over time, the United States developed civic institutions designed to manage industrial-scale private power. These institutions did not replace private enterprise, nor did they eliminate its capacity for innovation and growth. Instead, they introduced structure - frameworks through which large-scale systems could be guided, stabilized, and aligned with broader public needs.

Alongside these institutional developments, public utility systems were established to provide essential services - electricity, water, communication - on terms that extended beyond purely local or private

arrangements. These systems recognized that certain forms of infrastructure had become foundational to civic and economic life, requiring coordinated provision and broad access.

Then, in the face of economic collapse during the 1930s, the democratic challenge intensified. The Great Depression was visible, immediate, and widespread. Employment fell sharply. Economic activity contracted. Public confidence in existing systems weakened. The capacity of democratic institutions to respond effectively was tested under severe conditions.

The response again involved construction at scale. Through programs associated with the New Deal, the federal government undertook a wide range of public works. Roads, bridges, public buildings, and civic spaces were built across the country. These projects provided immediate employment, but they also created durable assets - structures that would support public life for decades to come.

At the same time, these efforts reinforced a broader principle: that democratic governance could act at scale to stabilize conditions when existing systems proved insufficient. Taken together, these developments addressed two closely related challenges. They introduced civic institutions capable of managing industrial-scale systems, and they built capacity to respond when those systems faltered.

The aim was not theoretical. It was practical: to maintain the conditions under which economic life could function and democratic participation could continue. As in earlier periods, these efforts were not without debate or difficulty. But over time, they contributed to a more stable relationship between large-scale economic activity and the public it affected. When power and instability challenged democracy, Americans built oversight and capacity.

Transition to the Modern Era

By the mid-twentieth century, the United States had developed systems to make governance visible, to maintain connection across distance, and to manage the scale and impact of industrial and economic life. But another shift was underway.

The pace of activity - economic, social, and technological - was increasing. Movement of people and goods accelerated. Communication began to extend beyond physical transport into electronic forms. Coordination across large systems became not only desirable, but necessary.

The next democratic challenge would not arise from distance alone, or from industrial scale alone. It would arise from the growing need to coordinate a complex society at speed. Once again, the question would be how to build systems capable of meeting that condition.

Part IV - Coordinating at Speed and Scale (Late 1900s-2000s)

Late 1900s to Early 2000s: Networks for a Fast-Moving Society

By the mid-twentieth century, the United States had constructed systems that made democratic life visible, connected a continental republic, and introduced structures to manage industrial scale and economic instability. A new democratic challenge was beginning to emerge.

The pace of modern life was accelerating. Economic activity moved more quickly across regions. Populations became more mobile. The coordination of goods, services, and public activity increasingly depended not only on connection, but on speed and reliability at scale. A society operating at this tempo required infrastructure capable of supporting continuous, large-scale coordination.

The response, once again, took the form of construction. The Federal-Aid Highway Act of 1956 initiated the development of the Interstate Highway System, creating an extensive network of roads designed for high-speed, long-distance travel. This system reshaped patterns of movement across the country, enabling goods and people to move with a consistency and efficiency not previously possible.

Beyond transportation, the Interstate system supported broader forms of coordination. It linked regional economies, facilitated national distribution systems, and allowed for more rapid response in times of need. It became part of the underlying structure through which a large and complex society could function with a greater degree of coherence.

At roughly the same time, another form of infrastructure began to take shape - less visible, but no less consequential. Early digital networks, developed through efforts such as ARPANET, introduced new possibilities for communication. Information could move independently of physical transport, transmitted electronically across connected systems. Over time, these networks expanded, forming the basis of what would become the modern internet.

This development marked a significant shift. For much of American history, the movement of information had been tied to physical systems - roads, railways, postal routes. With the emergence of

digital networks, communication began to operate within a different kind of environment: one defined by speed, scale, and continuous connectivity.

The effects were substantial. Barriers to communication were lowered. Individuals and institutions could exchange information rapidly across distance. New forms of civic participation became possible, as citizens gained access to a wider range of perspectives and channels for engagement.

At the same time, these systems introduced new dynamics. Information could circulate at a pace that exceeded traditional forms of verification and deliberation. The volume of available content increased dramatically. Digital platforms, many operated by private entities, began to play a central role in shaping how information was organized, presented, and encountered.

These developments did not replace earlier democratic challenges. They altered the conditions under which those challenges now unfolded. The infrastructure of connection had evolved into an infrastructure of continuous communication.

As in earlier eras, the construction of these systems addressed a real and pressing need: the ability to coordinate a large and fast-moving society. They succeeded in enabling forms of interaction and organization that would have been difficult to achieve otherwise. But they also marked the beginning of a new phase in the relationship between infrastructure and democratic life.

The systems that now shaped communication were not only channels. They were environments - spaces within which civic life increasingly took place. And unlike earlier forms of infrastructure, these environments were not always directly visible or easily understood as shared public systems. When speed and scale challenged democracy, Americans built networks.

By the early twenty-first century, these networks had become deeply embedded in daily life. Communication, information exchange, and public discourse increasingly occurred within digital systems. These systems connected individuals across vast distances in real time. They enabled rapid coordination, broad participation, and new forms of civic engagement.

At the same time, they introduced a different kind of condition. Civic life was no longer shaped only by visible institutions or physical infrastructure. It was increasingly influenced by a digital environment - dynamic, continuously evolving, and not always clearly seen as a whole. The challenge was no longer only how to connect or coordinate. It was how to understand the environment within which those connections and forms of coordination now occurred.

Part V - The 21st Century: Understanding the Civic Digital Field

The Present: A Challenge of Conditions

By this point, a pattern is visible. Across three centuries, when the conditions of democratic life changed, the United States responded by building infrastructure designed to meet those conditions. These responses differed in form - civic buildings, transportation networks, institutional frameworks, and digital systems - but they shared a common purpose: to sustain the practice of self-government under new circumstances.

The present moment introduces a different kind of democratic challenge. It is not defined by distance alone, or by industrial scale alone, or by the speed of physical coordination. It arises from the conditions of a civic environment that is increasingly digital, continuously active, and not always clearly seen as a whole.

Many citizens experience its effects without encountering it directly. Public discourse can feel more fragmented. Shared understanding can be more difficult to establish and maintain. Moments of civic disruption can emerge quickly, intensify, and then recede, often without clear resolution. At other times, changes occur more gradually, shaping expectations, perceptions, and patterns of interaction over longer periods.

These experiences do not point to a single event or cause. They suggest a shift in the conditions under which democratic life now unfolds. In earlier periods, key democratic challenges were often visible. Distance could be measured. Industrial systems could be observed. Economic collapse could be recognized in immediate terms.

The present challenge is different. The civic environment itself - within which information circulates, attention forms, and collective understanding develops - has become more complex and less directly observable. Its dynamics are distributed across systems, platforms, and interactions that do not present themselves as a single, shared public space.

At the same time, powerful digital systems are being deployed at speed, including increasingly influential AI systems operating in public and quasi-public environments. The risk is not only that crisis may emerge within this environment, but that crisis may arrive before the necessary democratic infrastructure for visibility, restraint, and public understanding has been built.

This introduces a particular kind of risk. A democracy depends on the ability of its citizens to develop sufficiently shared understandings of public life, to coordinate around common concerns, and to sustain a level of trust in the processes through which decisions are made. When the conditions supporting those capacities become difficult to see or assess, strain can develop - not always as a single, visible crisis, but as a series of smaller disruptions and gradual shifts that accumulate over time.

Some of these disruptions are visible. They can take the form of contested events, rapid cycles of information and reaction, or moments when public attention becomes sharply divided. Others are quieter. They may involve the steady fragmentation of shared reference points, the increasing difficulty of distinguishing signal from noise, or a gradual erosion in confidence that common ground can be established.

Both forms matter. Democratic life is shaped not only by moments of decision, but by the conditions that make those moments meaningful. When those conditions become unstable or difficult to interpret, the capacity for self-government can be affected - even if no single event fully captures the change.

This does not imply that earlier challenges have been resolved or replaced. It suggests that they now operate within a different environment - one that influences how they are experienced, understood, and addressed.

The pattern observed in earlier periods offers a point of reference. When democratic challenges have altered the conditions of civic life, durable responses have often involved building systems that made those conditions more visible, more structured, and more open to public understanding. The present moment raises a similar question.

If the civic environment in which democratic life unfolds has become more complex and less directly visible, what forms of infrastructure might help make its conditions more understandable? How might a democracy maintain its capacity for shared understanding, coordination, and trust within such an environment?

These questions do not point to a single answer. They suggest a space for inquiry, one that begins with recognizing the nature of the challenge itself.

Continuity and Responsibility

Across American history, democracy has not been maintained by assumption. It has been maintained through attention to changing conditions, and through the construction of systems capable of meeting them. The present moment does not stand outside that pattern. It extends it.

The conditions of democratic life have changed again - this time into a civic environment that is digital, dynamic, and not always easily seen. How that change is understood, and how it is addressed, remains an open question.

This series does not begin from alarm alone. It begins from architecture - asking whether, in the face of an emerging civic digital environment, the United States can once again build the kind of purpose-built infrastructure that allows democracy to survive a crisis not yet fully seen.

This series begins from a simple premise: that understanding the conditions of our civic environment is a necessary step in maintaining the practice of self-government under those conditions. Earlier generations built roads to connect a republic, institutions to manage industrial scale, and networks to enable communication at speed.

The present generation faces a different task. Not to replace those efforts, but to consider whether new forms of infrastructure may be required to understand the environment in which democratic life now unfolds. More specifically, it may need to design and build explicitly civic forms of digital infrastructure capable of illuminating the civic digital field before crisis forces the republic to confront these questions all at once.

That work has taken many forms across American history. This series will consider what American democracy may now be asking of us. Could a new civic form of public digital infrastructure, designed ahead of crisis, help stabilize our self-governance for generations to come?

Series Introduction:

Structured Complementarity
Civic AI Alignment Theory
and the
U.S. Non-Agentic Civic AI
Visibility Observatory System
Proposal

Structured Complementarity Civic AI Alignment Theory

STRUCTURED COMPLEMENTARITY CIVIC AI ALIGNMENT (SCCAI-A) recognizes existing civic capacities and their potential for complementarity while respecting the established roles of democratic and legal institutions and governmental processes. It insists that Civic AI alignment should shift from the corporate programming of machine values to the public arrangement of civic capacities. Specifically, this theory focuses on supporting the **Citizenry's Voluntary Distributed Civic Capacity**, **Continuous-Inquiry Interdisciplinary Expert Committees**, and **Non-Agentic Civic AI** (NACAI) as complementary civic capacities.

The alignment sought is not an alignment of machine values with civic values, nor an attempt to model, predict, or represent collective democratic conscience. It is a constitutional division of labor in which voluntary citizen authority, accountable continuous inquiry, technological illumination, and lawful institutional action retain distinct roles. Citizens require greater shared visibility into complex civic conditions; interdisciplinary committees require accessible, continuously organized evidence; and NACAI's observations acquire legitimate public meaning only through human conscience, democratic deliberation, and lawful institutional judgment. Together, these complementary civic capacities support civic problem-solving and civic data archiving.

Within this framework, NACAI serves as **Civic Visibility Public Infrastructure** - an auditable, non-agentic infrastructure that preserves public records, organizes and reveals relevant conditions and aggregate patterns, identifies system-level pressures, and helps institutions and communities recognize emerging conditions earlier. It must not score individuals, recommend policy outcomes, trigger enforcement, or exercise delegated governmental judgment. Judgment, political deliberation, and public responsibility remain with human citizens and institutions.

The theorist proposes a **U.S. Non-Agentic Civic AI Visibility Observatory System** within the nonpartisan, civilian **Library of Congress** as a potential institutional home for coordinating NACAI's civic visibility function. Its role must be archival as well as illuminating: preserving, authenticating, and making intelligible the evidentiary record in a public environment where civic data is reproduced, compressed, recombined, and interpreted at scale. Without provenance-based authentication - the verification of civic data's origin and integrity - democratic judgment cannot remain reliably grounded in evidence. Core theoretical principles, presented in *Non-Agentic Civic AI - Illuminating America's Civic Digital Field*, include No Moral Outsourcing, Digital Nature Integrity, Amoral AI Civic Advantage, and Moral Visibility.

Structured Complementarity Test: A Civic AI arrangement is acceptable only if it strengthens the differentiated and mutually supportive capacities of citizens, public institutions, and bounded computational systems without making those capacities interchangeable or obscuring human responsibility.

Structured Complementarity rejects the assumptions that Civic AI is inherently either dangerous or benign; its civic admissibility depends upon its design, bounded institutional role, operational limits, accountability structure - and above all, its conformity with enduring democratic principles, as proposed in the **52-Gates Democratic Permissions Certification Threshold.** It also rejects the premise that technological optimization should attain governing influence. Deployment in government environments shall not logically confer Civic AI status. Non-Agentic Civic AI may offer democracies high capability without becoming unaccountable civic power or eroding democratic self-government.

Structured Complementarity Civic AI Alignment Theory

Introduced in U.S. Civic AI Field Context

- Alignment Traditions and the Missing Civic Layer -

(From the Preface of *Non-Agentic Civic AI – Illuminating America's Civic Digital Field*)

Christopher Hunt Robertson, M.Ed.

With the Collaborative Support of Advanced AI – ChatGPT, Perplexity, and Gemini/Google

Democracies cannot govern what they cannot see.

Over the past decade, AI alignment has become a wide-ranging and urgent conversation. Some thinkers focus on hypothetical future superintelligences and existential risk. Others document present harms in welfare systems, labor markets, and digital platforms. Still others examine how attention, autonomy, and democratic norms erode under commercial and technological pressure.

Each of these traditions illuminates something essential, and none should be ignored.

This book series' contribution, ***Structured Complementarity Civic AI Alignment,*** stands alongside those efforts, but it does not attempt to replicate them. Its focus is narrower and more concrete: the civic conditions under which any alignment work - technical, regulatory, or philosophical - can succeed in a democracy.

(Beginning of Theoretical Statement)

STRUCTURED COMPLEMENTARITY CIVIC AI ALIGNMENT (SCCAI-A) recognizes existing civic capacities and their potential for complementarity while respecting the established roles of democratic and legal institutions and governmental processes. It insists that Civic AI alignment should shift from the corporate programming of machine values to the public arrangement of civic capacities. Specifically, this theory focuses on supporting the **Citizenry's Voluntary Distributed Civic Capacity**, **Continuous-Inquiry Interdisciplinary Expert Committees**, and **Non-Agentic Civic AI** (NACAI) as complementary civic capacities.

The alignment sought is not an alignment of machine values with civic values, nor an attempt to model, predict, or represent collective democratic conscience. It is a constitutional division of labor in which voluntary citizen authority, accountable continuous inquiry, technological illumination, and lawful institutional action retain distinct roles. Citizens require greater shared visibility into complex civic conditions; interdisciplinary committees require accessible, continuously organized evidence; and NACAI's observations acquire legitimate public meaning only through human conscience, democratic deliberation, and lawful institutional judgment. Together, these complementary civic capacities support civic problem-solving and civic data archiving.

Within this framework, NACAI serves as **Civic Visibility Public Infrastructure** - an auditable, non-agentic infrastructure that preserves public records, organizes and reveals relevant conditions and aggregate patterns, identifies system-level pressures, and helps institutions and communities recognize emerging conditions earlier. It must not score individuals, recommend policy outcomes, trigger enforcement, or exercise delegated governmental judgment. Judgment, political deliberation, and public responsibility remain with human citizens and institutions.

The theorist proposes a **U.S. Franklin-Dewey-Citizenry Civic Data Observatory** within the nonpartisan, civilian **Library of Congress** as a potential institutional home for coordinating NACAI's civic visibility function. Its role must be archival as well as illuminating: preserving, authenticating, and making intelligible the evidentiary record in a public environment where civic data is reproduced, compressed, recombined, and interpreted at scale. Without provenance-based authentication - the verification of civic data's origin and integrity - democratic judgment cannot remain reliably grounded in evidence. Core theoretical principles, presented in *Non-Agentic Civic AI - Illuminating America's Civic Digital Field*, include No Moral Outsourcing, Digital Nature Integrity, Amoral AI Civic Advantage, and Moral Visibility.

Structured Complementarity Test: A Civic AI arrangement is acceptable only if it strengthens the differentiated and mutually supportive capacities of citizens, public institutions, and bounded computational systems without making those capacities interchangeable or obscuring human responsibility.

Structured Complementarity rejects the assumptions that Civic AI is inherently either dangerous or benign; its civic admissibility depends upon its design, bounded institutional role, operational limits, accountability structure - and above all, its conformity with enduring democratic principles, as proposed in the **52-Gates Democratic Permissions Certification Threshold**. It also rejects the premise that technological optimization should attain governing influence. Deployment in government environments shall not logically confer Civic AI status.

Non-Agentic Civic AI may offer democracies high capability without becoming unaccountable civic power or eroding democratic self-government.

(End of Theoretical Statement)

This constitutional division of labor is Deweyan in spirit. John Dewey understood democracy not simply as a system for selecting officials, but as a form of associated life sustained by communication, inquiry, and the continuing capacity of a public to recognize and address the consequences of collective action. Democratic intelligence, in this sense, is not the intelligence of an expert system acting on behalf of passive citizens. It is the shared human practice of bringing consequences into view, testing interpretations against experience, hearing those affected, revising provisional understandings, and deciding openly what ought to be done.

Non-Agentic Civic AI can assist that practice only on its proper side: it may help make conditions more observable, records more preservable, patterns more comparable, and uncertainty more explicit. It may not stand in for the public inquiry through which people determine the meaning of those conditions, deliberate over competing obligations, and authorize lawful action. The objective is not automated collective intelligence, but better democratic conditions for human collective intelligence.

It begins from a simple claim: if free societies cannot see clearly what they are doing at the speeds and scales now shaped by AI, no technical solution can keep them aligned. Moral visibility - the capacity to perceive emerging civic harms, uncertainty, accountability pathways, and system-level drift in time to deliberate - is not a luxury. It is infrastructure.

What follows briefly sketches four influential alignment and AI-ethics traditions, what each makes visible, and the civic gap this work is designed to address.

Long-Termist Alignment: Nick Bostrom and Existential Risk

Nick Bostrom's *Superintelligence* crystallized a powerful concern: future machine intelligences may outstrip human control and pose existential danger. His work operates on long time horizons and at a planetary scale, emphasizing technical control problems, misaligned objectives, and the need for coordination around advanced systems.

This book does not contest those risks. It asks a prior civic question. Before humanity worries about steering hypothetical superintelligences, can existing institutions reliably see what today's systems are already doing? Before debating the goals of future agents, can democracies track hazards created by present systems operating at machine speed in schools, hospitals, courts, media systems, and elections? Where long-termist alignment looks far forward, this work focuses on a prerequisite layer: the visibility and responsibility conditions within the societies alignment is meant to protect. If those conditions collapse, even excellent technical alignment work may arrive too late, or land inside institutions too blind to govern it well.

In this sense, AI appears here not first as an existential adversary, but as an amplifier of already-present civic weaknesses: fragmented shared reality, diminishing democratic visibility, and institutions required to decide under conditions they can no longer adequately perceive.

Justice-First Analysis: Virginia Eubanks and Structural Harm

Virginia Eubanks's *Automating Inequality* documents how data systems and automation can deepen inequality in housing, welfare, and public services. Her work is empirically grounded and morally clear, showing how opaque systems can quietly punish poor and marginalized people. It remains a vital account of structural harm.

This work shares that concern for human dignity, but extends the analysis from diagnosis to civic design. It treats moral visibility as a democratic requirement: the ability of institutions and publics to see where harm is forming, who is bearing it, how patterns propagate across systems, and who retains responsibility for a response.

Where Eubanks shows how current systems can fail vulnerable communities, this work asks what kinds of non-agentic civic instruments might help repair that failure: instruments that monitor systems rather than individuals; surface hidden harms, uncertainty, and overlooked stakeholders without judging or scoring anyone; and reveal structural risk before it hardens into entrenched injustice.

Eubanks diagnoses social harm. This book proposes public-works responses: civic infrastructure designed to strengthen responsibility, public understanding, and democratic visibility without replacing human judgment.

Surveillance Capitalism Critique: Shoshana Zuboff and Autonomy

Shoshana Zuboff's *The Age of Surveillance Capitalism* offers a sweeping critique of how digital platforms extract behavioral data and commodify prediction. Her analysis portrays the fusion of capital and computation as a regime that threatens human autonomy by treating lived experience as raw material for prediction and influence.

This work shares Zuboff's concern that people are increasingly treated as inventory. It diverges in emphasis. Rather than centering economic extraction alone, it focuses on the epistemic and civic consequences: how citizens, officials, and institutions perceive risk in an environment where manipulation, distortion, and automated amplification reshape public life faster than democratic processes can respond.

Where Zuboff asks who owns behavioral data and the futures it enables, this work asks an adjacent question: how can democracies regain enough shared sight to govern responsibly in that environment? Its answer is a civic blueprint rather than a political-economy critique: a non-agentic, rights-preserving ensemble of interpretive tools that illuminates hazards without centralizing power, scoring citizens, or directing action.

Attention and Design Ethics: James Williams and Moral Focus

James Williams's *Stand Out of Our Light* examines the attention economy and argues that persuasive technologies erode conditions necessary for individual freedom. His work is centered on personal agency, the capacity to attend to what one finds meaningful.

This book shares Williams's concern for moral attention and inner freedom, but situates it within a Benjamin Franklin–inspired ethic of public works. Individual conscience remains the non-delegable source of moral judgment. Under contemporary conditions, private attention alone is not enough. The systems surrounding conscience now operate at speeds and scales that can overwhelm even the most disciplined individual.

Where Williams calls for more humane design at the level of products and interfaces, this work seeks to institutionalize clarity as a public good. It imagines civic-scale institutions able to say: "Here is what we saw. Here is what we do not know. Here is what we refused. Here is who owns the decision."

In this way, expressions of conscience are given something increasingly rare in digital life: a stable place to stand.

The Missing Civic Layer

Taken together, these traditions illuminate technical control problems, structural injustice, economic extraction, and the erosion of attention. Each addresses an essential facet of the alignment challenge. Yet across their differences, one operating layer is rarely treated as a primary design target: the civic infrastructure through which democratic societies perceive, interpret, contest, and respond to emerging harms.

None of these approaches is principally designed to answer questions such as:

- How can public institutions remain functional when hazardous digital patterns form faster than human perception?
- How can dispersed consequences become matters that citizens and communities can recognize, investigate, and deliberate about together?
- How can democracies restore shared visibility without automating moral judgment or centralizing coercive authority?
- How might non-agentic tools strengthen human judgment, preserve the right to refuse, and make responsibility visible rather than diffuse it?
- How can a society develop greater civic visibility without making citizens more legible to power?

That final question requires particular care. In *Seeing Like a State*, James C. Scott warned that institutions may simplify complex social realities in order to make populations administratively legible, discounting local knowledge and enabling overconfident intervention from above. His warning is

not an argument against public knowledge. It is a warning against confusing an administrative representation of society with the society itself, and against allowing simplified visibility to become an instrument of centralized control.

The purpose of the civic visibility proposed here is therefore not to make citizens more visible to power. It is to make consequential institutional patterns, public risks, uncertainty, and lines of responsibility more visible to citizens, communities, and accountable public servants. Its instruments must remain rights-bounded, publicly auditable, explicit about what they cannot know, and open to correction by people whose lived knowledge reveals what any formal system has missed.

In the language of this work, the missing layer is ***Civic Capacities Alignment,*** the constitutional principle that AI's civic capacities must remain aligned with the limited role democracy permits them to perform. Its practical expression is ***Civic Visibility Public Infrastructure*** **-** rights-bounded public institutions, procedures, data practices, and non-agentic instruments through which a society can develop and sustain moral visibility.

Moral visibility is the practical capacity of a free society to perceive consequential patterns, emerging harms, uncertainty, accountability pathways, and system-level drift in time to deliberate. It joins a Benjamin Franklin-inspired public-works ethic to a John Dewey-inspired democratic ethic of inquiry: institutions should enlarge the public's capacity to learn from shared consequences without transferring moral or political authority to technical systems. Franklin supplies the practical imperative to build durable civic instruments; Dewey supplies the democratic requirement that those instruments remain communicative, publicly intelligible, corrigible, and answerable to the people whose common life they illuminate.

This book sketches one possible expression of that infrastructure: a ***Franklin - Dewey - Citizenry Civic Data Observatory System***, built on two fixed commitments - human moral primacy and non-agentic design.

Its ambition is modest but durable. It seeks to extend perception rather than automate control; to make risks, uncertainty, and responsibility more legible without scoring citizens, directing conduct, or concentrating authority. Its governing division of labor is simple:

AI illuminates. Humans judge. Lawful institutions decide.

Like Franklin's public works, and consistent with Dewey's experimental democracy, such an architecture must be tested in the open: piloted in limited settings, publicly explained, independently audited, challenged by affected communities, revised in light of experience, and, where necessary, stopped.

This framework does not replace existing alignment work. It addresses a condition under which technical, regulatory, and philosophical alignment efforts can become democratically usable: societies must retain enough shared visibility to govern themselves.

Non-Agentic Civic AI is offered to fellow alignment thinkers not as a competing school, but as a practical invitation to widen alignment's design target - from models and incentives alone to the civic infrastructure that surrounds them. It invites exploration of how rights-preserving, non-agentic

instruments might extend early shared visibility across institutions during future public emergencies and other fast-moving civic crises, without automating judgment or centralizing power - giving humans a clearer map of risk while leaving the hardest choices in their hands.

This work does not claim to solve alignment. It proposes a public-works framework designed for open testing by any alignment agenda that takes democratic self-government seriously.

Civic Visibility Public Infrastructure must not make citizens legible to power; it must make power, risk, uncertainty, and institutional responsibility legible to the citizenry.

People must be able to collectively see civic digital signal disturbances early enough to make timely, informed decisions that may lessen harmful societal impacts and mitigate costly damages.

In the United States, federal, state, and local governments already commit vast public resources to recovering from digital disruptions in the public sphere, costs borne by taxpayers and the communities they serve. Prevention cannot eliminate every danger. But earlier shared visibility can help democratic institutions recognize consequential patterns, investigate them openly, and act before preventable harms cascade.

As Benjamin Franklin might remind us, "An ounce of prevention is worth a pound of cure." Let us build the ***Civic Visibility Public Infrastructure*** now required in the Age of Humanity with AI - not to make citizens more legible to power, but to ensure that people can see power, risk, uncertainty, and institutional responsibility clearly enough, soon enough, to sustain self-government.

Ben Franklin's Public Lighting

Whale Oil Street Lamp (American)

Ben Franklin significantly improved the street lamp's design and promoted its adoption in Philadelphia around 1757, introducing a funnel and vents to reduce smoke buildup in enclosed globes, making them cleaner and more efficient, and encouraging citizens to install them to light their streets. He proposed a design with four flat panes, a top funnel for smoke, and bottom vents for air, inspired by an idea from a neighbor, John Clifton, and used two wicks for brighter light. Franklin's streetlamp was designed to stay cleaner, break less easily, and cost less to maintain than the imported globe lamps Philadelphia had used before—an early example of American frugality in public works.

Ben Franklin Tercentenary (1706 – 2006)
Atwater Kent Museum of Philadelphia,
Historical Society of Pennsylvania Collection

(Photo Credit: Peter Harholdt)

The "Civic Digital Visibility First" Architectural Stack

- A New Civic AI Approach -

"Democracies cannot govern what they cannot see."

The **Civic Digital Visibility First Architectural Stack** treats civic digital visibility as the prerequisite condition of democratic governability. Rather than centering optimization, prediction, or automated decision authority, it proposes visibility-only public-works that restore shared sight and deliberative time under machine-speed conditions.

The stack's components - each within current technological reach - address civic digital invisibility, not domestic privacy regulation or citizen surveillance. Inspired by Benjamin Franklin's public-works ethic, the design is non-agentic, rights-tethered, revocable, and refusal-bound.

The Problem

Six Rings of Digital Invisibility and Civic Destabilization (Problem Diagnosis): The Six Rings provide the diagnostic lens for the entire architecture. They describe how fragmented information flows, opaque digital systems, and accelerating machine-speed processes erode shared civic visibility before overt constitutional crisis appears. By tracing how civic harms accumulate across successive "rings" of invisibility - from overload and attention capture to institutional blindness and destabilization - the model reframes alignment as a problem of shared perception and consequence representation rather than abstract technical control. The remaining components of the stack are designed to restore sight across these rings - without transferring authority from humans to machines.

The Proposed Architectural Solution

The U.S. Non-Agentic Civic AI Visibility Observatory System (Border-Facing National Visibility Infrastructure): The proposed Observatory is the border-facing visibility instrument of the stack. It is a purpose-built, non-agentic civic infrastructure designed to detect foreign-origin digital interference at aggregate and system levels. Modeled on Franklin's public works - lamps, levees, and lighthouses - the Observatory extends shared sight across institutions while leaving judgment and coercive authority entirely in human hands. It does not conduct domestic surveillance. It does not support enforcement. It does not classify persons. It detects foreign-origin systemic patterns and refers findings to existing constitutional authorities, who remain free to act - or not - with eyes open. The Observatory is admissible only under the 52-Gates Democratic Permissions Threshold and remains revocable, suspendable, and dismantlable by democratic decision.

The Architecture's Democratic Foundation

52-Gates Democratic Permissions Threshold (Constitutional Admissibility for All Civic AI): The 52-Gates of Democratic Permissions is a pass/fail constitutional instrument for admitting artificial intelligence systems near democratic authority. Unlike flexible, risk-based governance frameworks, the 52-Gates ask a prior question: under what structural conditions, if any, may machine systems operate near democratic power at all? Each Gate articulates a democratic principle, a structural necessity, and a binary disqualifying test. Systems that fail any gate are inadmissible. The 52-Gates instrument prohibits delegation of moral judgment to machines, forbids enforcement or optimization authority, requires public visibility of civic-scale outputs, and embeds refusal, revocability, and periodic re-certification as ordinary safeguards.
No system operating near democratic authority may qualify as constitutionally legitimate without passing every gate.

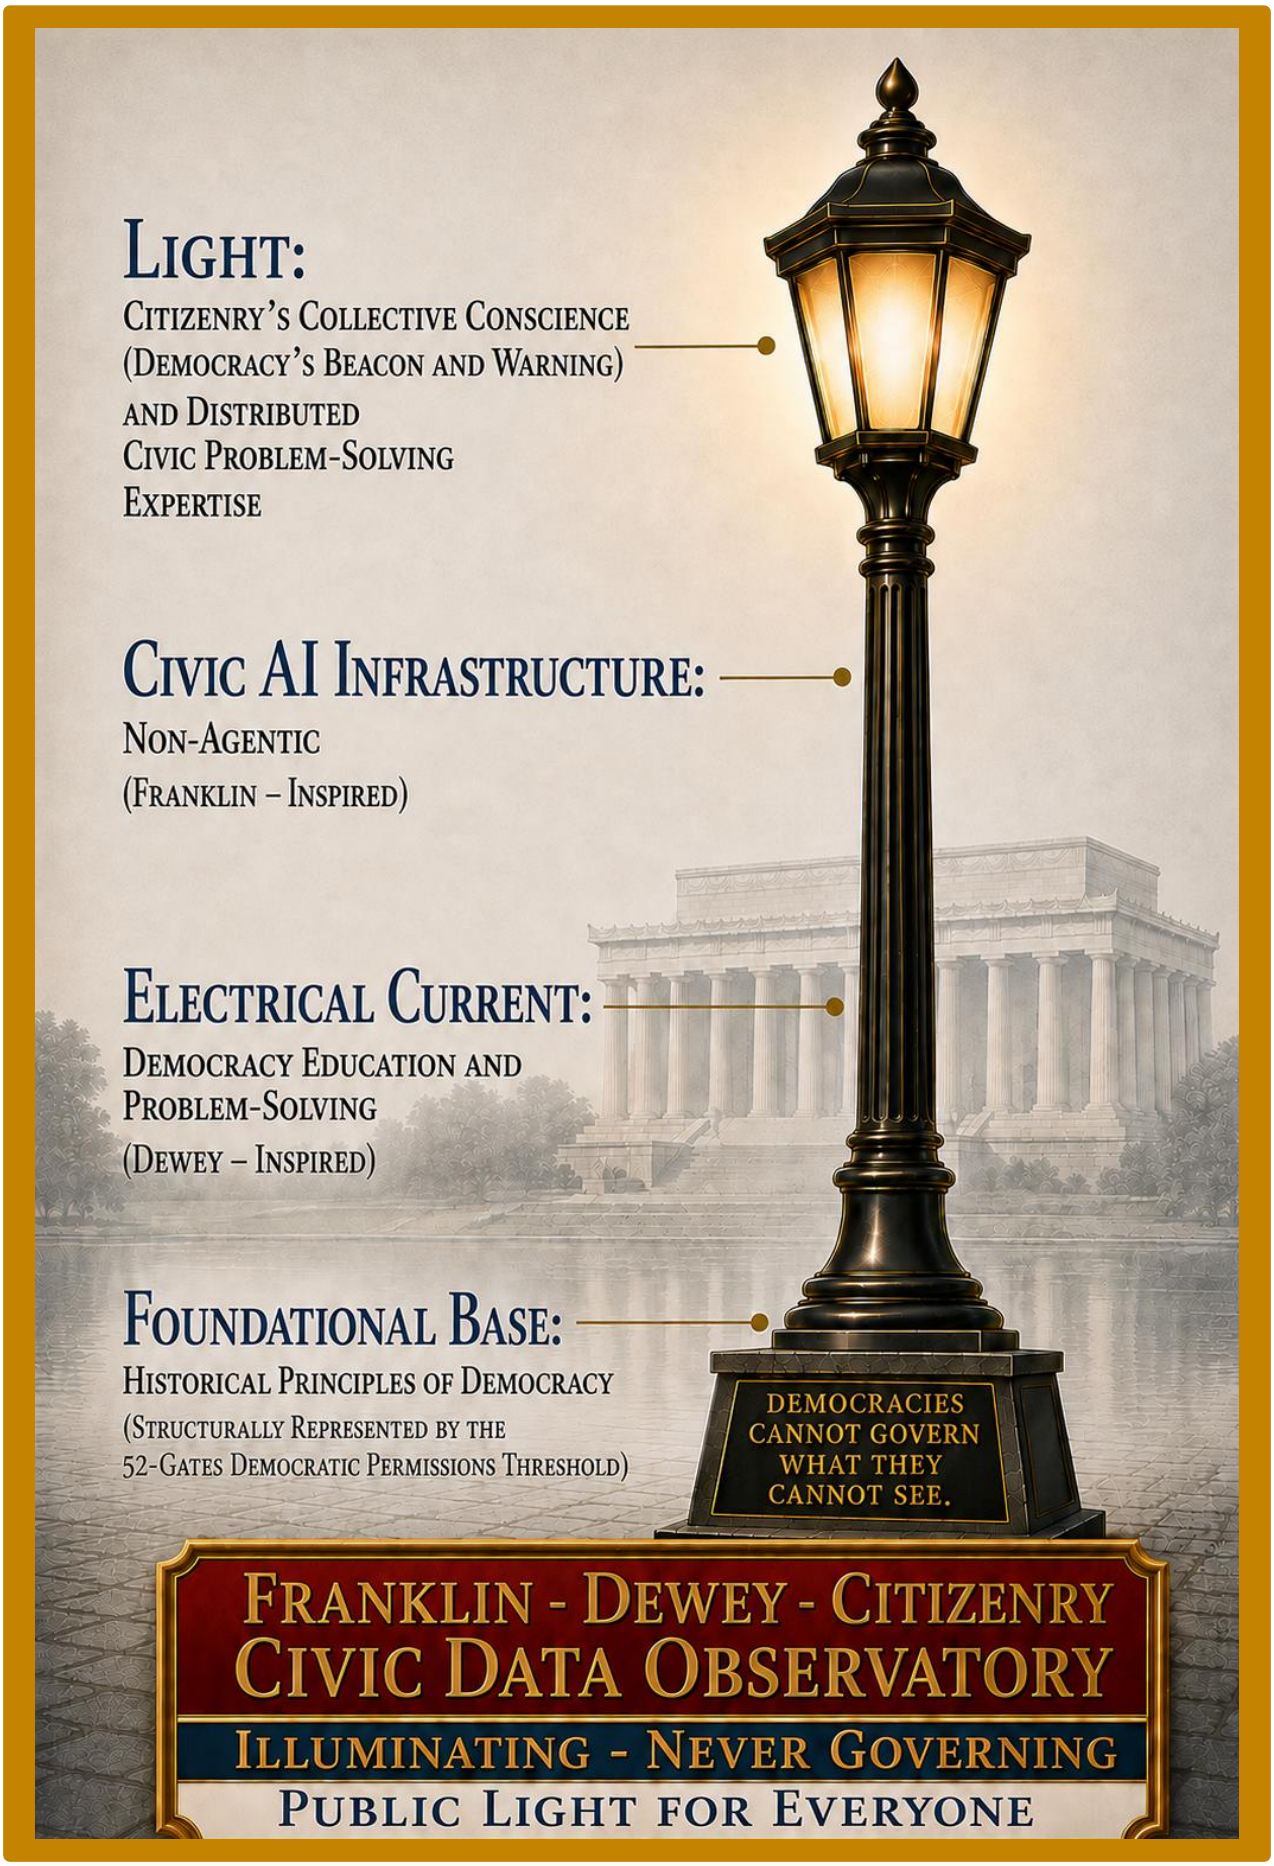

(Image by ChatGPT)

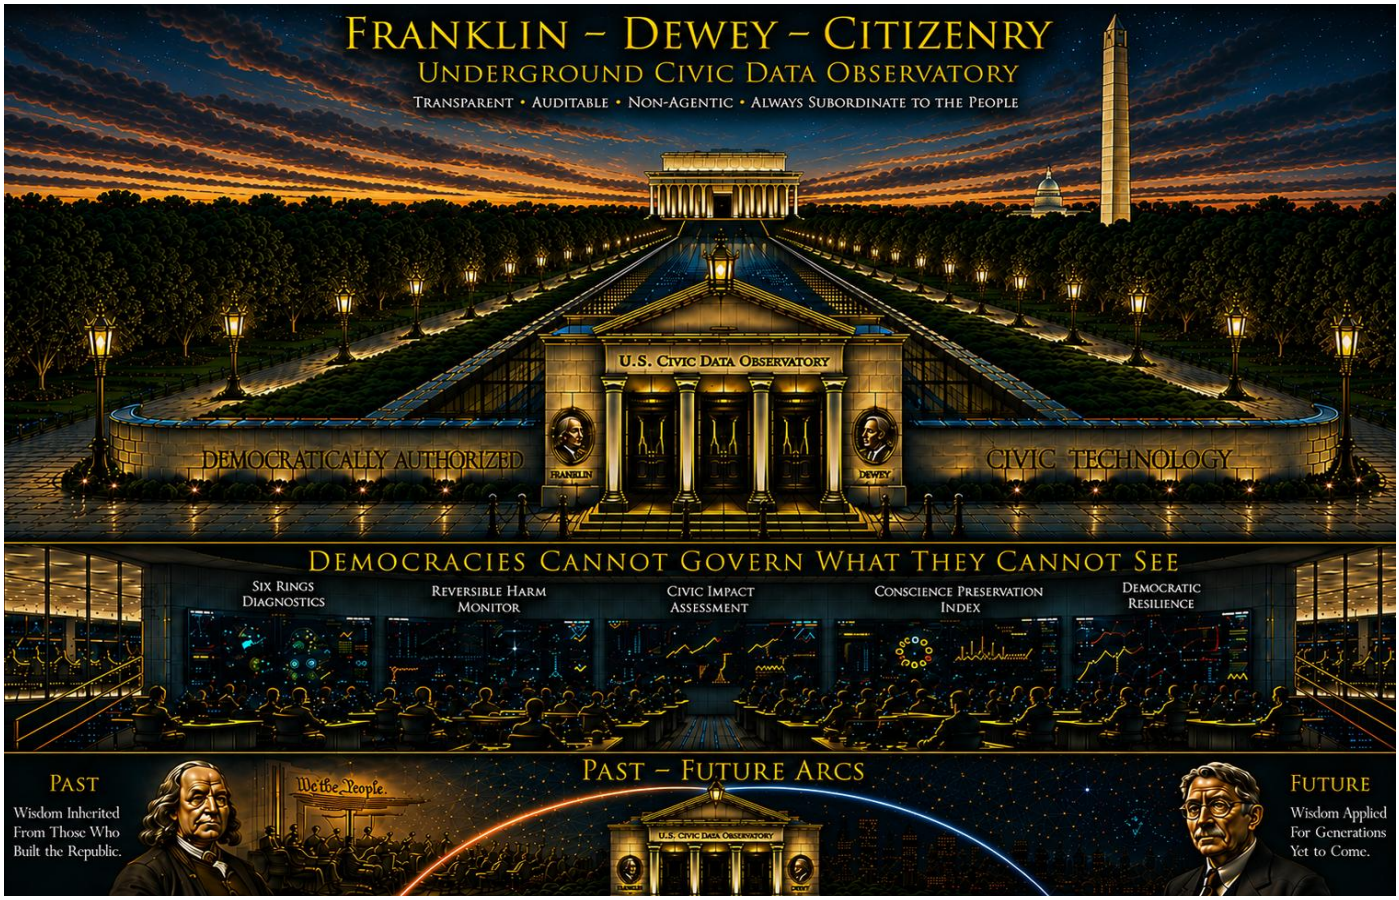
FRANKLIN – DEWEY – CITIZENRY
UNDERGROUND CIVIC DATA OBSERVATORY
TRANSPARENT • AUDITABLE • NON-AGENTIC • ALWAYS SUBORDINATE TO THE PEOPLE
U.S. CIVIC DATA OBSERVATORY
FRANKLIN
DEWEY
DEMOCRATICALLY AUTHORIZED
CIVIC TECHNOLOGY
DEMOCRACIES CANNOT GOVERN WHAT THEY CANNOT SEE
SIX RINGS DIAGNOSTICS
REVERSIBLE HARM MONITOR
CIVIC IMPACT ASSESSMENT
CONSCIENCE PRESERVATION INDEX
DEMOCRATIC RESILIENCE
PAST
Wisdom Inherited From Those Who Built the Republic.
We the People.
PAST – FUTURE ARCS
U.S. CIVIC DATA OBSERVATORY
FUTURE
Wisdom Applied For Generations Yet to Come.

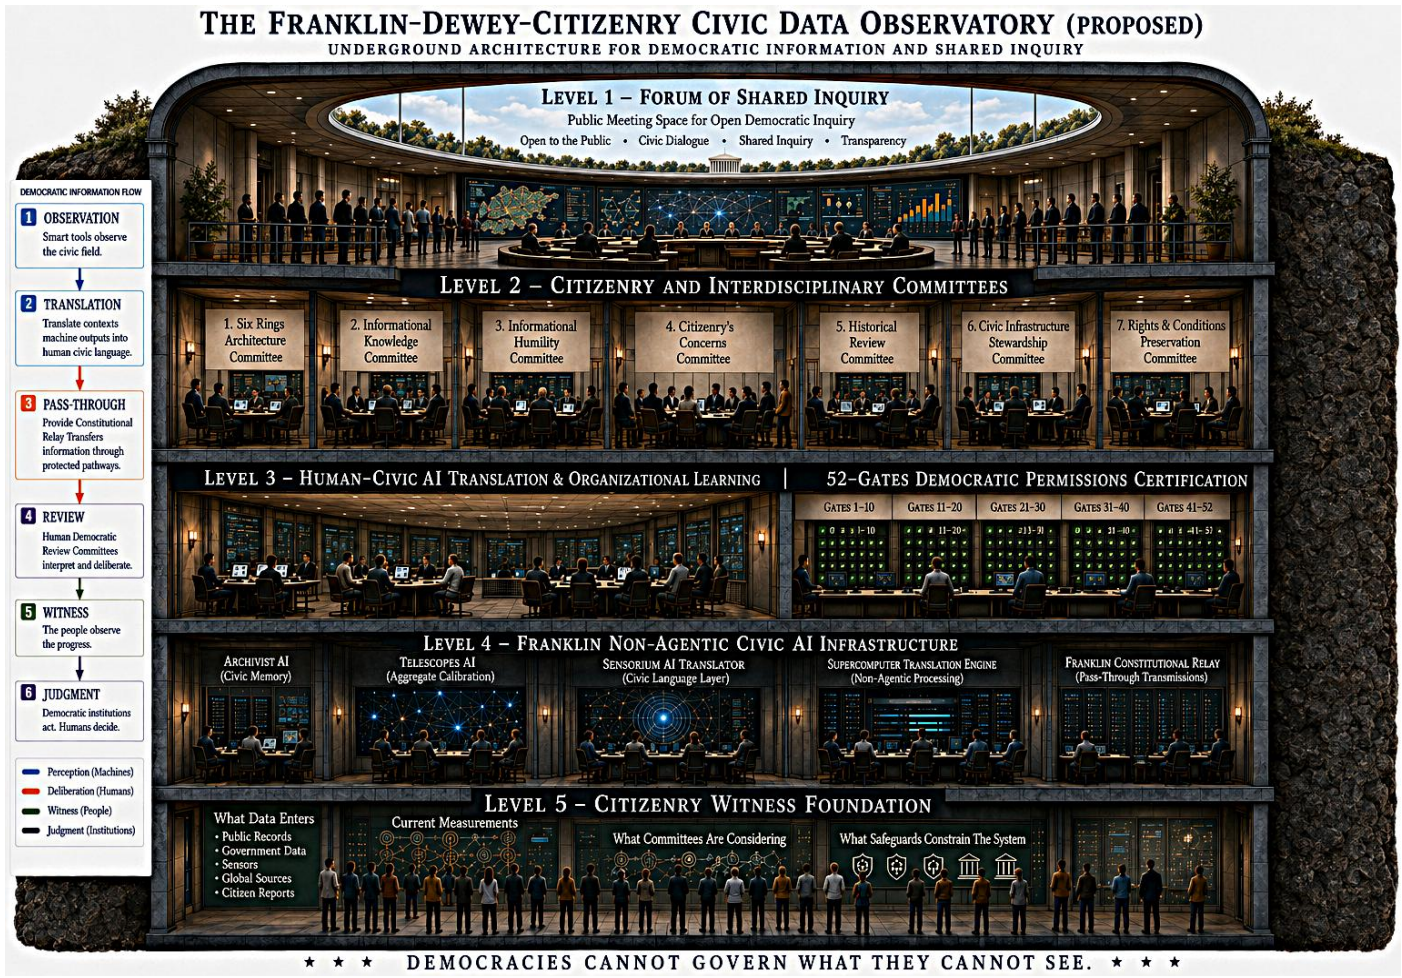
THE FRANKLIN-DEWEY-CITIZENRY CIVIC DATA OBSERVATORY (PROPOSED)
UNDERGROUND ARCHITECTURE FOR DEMOCRATIC INFORMATION AND SHARED INQUIRY
DEMOCRATIC INFORMATION FLOW
1 OBSERVATION
Smart tools observe the civic field.
2 TRANSLATION
Translate contexts machine outputs into human civic language.
3 PASS-THROUGH
Provide Constitutional Relay Transfers information through protected pathways.
4 REVIEW
Human Democratic Review Committees interpret and deliberate.
5 WITNESS
The people observe the progress.
6 JUDGMENT
Democratic institutions act. Humans decide.
Perception (Machines)
Deliberation (Humans)
Witness (People)
Judgment (Institutions)
LEVEL 1 – FORUM OF SHARED INQUIRY
Public Meeting Space for Open Democratic Inquiry
Open to the Public • Civic Dialogue • Shared Inquiry • Transparency
LEVEL 2 – CITIZENRY AND INTERDISCIPLINARY COMMITTEES
1. Six Rings Architecture Committee
2. Informational Knowledge Committee
3. Informational Humility Committee
4. Citizenry's Concerns Committee
5. Historical Review Committee
6. Civic Infrastructure Stewardship Committee
7. Rights & Conditions Preservation Committee
LEVEL 3 – HUMAN-CIVIC AI TRANSLATION & ORGANIZATIONAL LEARNING
52-GATES DEMOCRATIC PERMISSIONS CERTIFICATION
GATES 1–10
GATES 11–20
GATES 21–30
GATES 31–40
GATES 41–52
LEVEL 4 – FRANKLIN NON-AGENTIC CIVIC AI INFRASTRUCTURE
ARCHIVIST AI (Civic Memory)
TELESCOPES AI (Aggregate Calibration)
SENSORIUM AI TRANSLATOR (Civic Language Layer)
SUPERCOMPUTER TRANSLATION ENGINE (Non-Agentic Processing)
FRANKLIN CONSTITUTIONAL RELAY (Pass-Through Transmissions)
LEVEL 5 – CITIZENRY WITNESS FOUNDATION
What Data Enters
• Public Records
• Government Data
• Sensors
• Global Sources
• Citizen Reports
Current Measurements
What Committees Are Considering
What Safeguards Constrain The System
★ ★ ★ DEMOCRACIES CANNOT GOVERN WHAT THEY CANNOT SEE. ★ ★ ★

Preparing
for
Civic Digital Storms Ahead

The Peril of the Unanchored Machine: Why Precedent is Infrastructure

We are deploying advanced artificial intelligence across the civic landscape with a dangerous form of historical amnesia.

As public agencies rush to adopt autonomous agents, optimization frameworks, and predictive modeling, the prevailing conversation treats AI alignment as an entirely novel, purely technical control problem. We focus intensely on corporate machine values or hypothetical future superintelligences. In doing so, we fail to recognize that the most pressing challenge of our era is not technical—it is civic.

By treating the digital field as an unprecedented wilderness, we ignore a deep, uniquely American lineage of intellectual, infrastructural, and physical models specifically built to protect self-governance from destabilizing forces. We bypass these proven models at our own profound civilizational peril.

A democracy cannot long endure when its citizens can no longer see the forces shaping their collective life, or when the speed of institutional processes completely outpaces human deliberative time. When the conditions of self-government shift, the answer is never to surrender human moral authority to an unaccountable computational engine.

The answer is to build public infrastructure that restores shared sight.

Agentic AI asks democracy to accommodate a machine. The models that follow demand that our machines accommodate democracy.

Let us look back, clearly and urgently, before we attempt to build forward.

Franklin and Dewey: Inspiring an Infrastructure Ethic of Analysis and Visibility

An Intellectual Foundation for a New Observatory-and-Lighthouse for the Digital Age

Benjamin Franklin and John Dewey offer complementary starting points for a democratic question: How can a community understand its shared circumstances well enough to improve them together? Their work inspires our contemporary synthesis: an infrastructure ethic of analysis and visibility in service of democratic inquiry and self-government. Public-serving systems should help citizens see, understand, question, and address matters of common concern. Visibility is the beginning—not the completion—of democratic responsibility.

Franklin joined practical ingenuity with civic cooperation. His contributions to libraries, firefighting associations, public communication, and other shared institutions demonstrated how organized knowledge and voluntary participation could strengthen the common good. He inspires our practical imperative: build useful things, make their benefits shareable, and enable people to participate meaningfully in their support and improvement.

Dewey explains why useful institutions also require effective communication and continuing inquiry. In *The Public and Its Problems*, he examined the difficulty of recognizing and organizing around the far-reaching consequences of an increasingly complex society. Citizens cannot adequately address common problems whose consequences remain obscure. Expertise matters, but information alone does not constitute democratic understanding. Citizens need opportunities to examine evidence, question explanations, assess consequences, and judge how expert knowledge bears upon their shared lives. Dewey inspires our inquiry imperative: make accessible information a resource for investigation, deliberation, and responsible collective judgment.

The Digital Observatory-and-Lighthouse proposed herein gives institutional expression to these complementary commitments. The Observatory preserves records and examines evidence; the Lighthouse communicates findings in accessible forms. The Lighthouse also signals emerging pressures and provides evidence-grounded warnings for public attention and human judgment. Non-agentic civic AI assists these observational, analytical, and communicative functions. All such tools serve informational—not governing—roles. Accountable human institutional gatekeepers oversee access and public release under defined rules and appropriate review. Privacy, lawful confidentiality, and security remain essential boundaries.

Franklin supplies the impulse to build for common benefit; Dewey supplies the discipline of shared inquiry. Our synthesis places both in service of democratic infrastructure that makes the conditions of common life more intelligible without transferring governing authority from people to machines. Its purpose is not indiscriminate disclosure or automated government, but public understanding that strengthens citizens' capacity to investigate, deliberate, judge, and act.

The infrastructure helps illuminate the questions. The citizenry retains the authority to ask them, deliberate upon them, and decide what should follow.

Structurally Protecting American Democracy in the Digital Age:

The Proposed U.S. Non-Agentic Civic AI Visibility Observatory System

(Gated Observatory-and-Lighthouse Public Infrastructure, with Human Institutional Gatekeepers and the Citizenry-at-the-Helm)

America faces a deluge of civic data, both seen and unseen. Yet the accumulation of information does not ensure that a democracy can understand the conditions confronting it. Civic data must be examined, consequential anomalies identified, and emerging hazards made visible so that citizens and accountable institutions can deliberate and act with informed judgment.

The Observatory and the Lighthouse represent complementary functions of the proposed civic infrastructure. The Observatory examines. It brings evidence into view, preserves context, and reveals patterns that might otherwise remain hidden within the growing complexity of public life. The Lighthouse illuminates consequential findings, signaling anomalies, uncertainties, and potential hazards that warrant human attention.

Neither function is sufficient alone. Analysis without public illumination may leave consequential knowledge unnoticed or unused. A signal without adequate analysis risks producing alarm without sufficient grounding. Their integration connects the examination of civic conditions with the public recognition of findings that merit attention.

But visibility is not the destination. Any consequential action informed by the infrastructure's observations remains a matter for human deliberation and accountable institutional decision. The gated perimeter establishes enforceable limits around the infrastructure and defines permitted exchanges. It prohibits Agentic AI from entering or interacting with the protected system and prohibits the system itself from exercising delegated governmental judgment. Human institutional gatekeepers safeguard these boundaries under publicly established rules. Their responsibility is to administer legitimate access, protect rights, and preserve faithful disclosure—not to waive foundational prohibitions or claim ownership of civic visibility.

The proposed Non-Agentic Civic AI system can illuminate civic conditions and strengthen public inquiry without assuming the authority to govern. Citizens and accountable democratic institutions retain the responsibility to interpret evidence, weigh competing values, exercise conscience, and determine their course. The infrastructure informs their judgment; it does not replace it.

This is the purpose of integrating civic observation, public illumination, rights-bounded access, and accountable human institutional stewardship within a common infrastructure: not to automate democracy's moral responsibilities, but to strengthen the conditions under which those responsibilities can be exercised.

The Observatory examines. The Lighthouse illuminates. Human institutional gatekeepers safeguard the boundaries. The U.S. Citizenry retains the helm.

This model extends a conditional welcome to Non-Agentic Civic AI. May our democracy balance innovation with caution, drawing upon the technologies we create while preserving what we cannot afford to lose.

American Observatories and Lighthouses,
Inspiring Public Analysis and Illumination Infrastructure for Our Digital Age

The Toledo Harbor Lighthouse, built in 1903, continues to serve a vital dual purpose as a beacon of illumination and a crucial post for weather observation. For over six decades, dedicated keepers hand-operated its signature Fresnel lens—notable for its distinct internal red flash—to guide heavy freighter traffic, while simultaneously recording meticulous, manual logbook entries of local wind speeds, visibility, and severe Great Lakes storms.

Today, the historic brick lighthouse has successfully transitioned into a high-tech maritime station, seamlessly combining modern illumination with advanced, real-time meteorological reporting. Its modern solar-powered LED optic continues to project a vital safety beacon across Maumee Bay, while the structure itself serves as a host platform for NOAA's C-MAN Station THL01. Maintained by the Great Lakes Environmental Research Laboratory, this automated equipment continuously transmits critical public data on wind gusts, air temperature, and barometric pressure, complemented by a webcam.

Non-Agentic Civic A.I.: Illuminating America's Civic Digital Field (Vol. I)

CIVIC DIGITAL VISIBILITY FIRST

- Part I -

SIX RINGS of DIGITAL INVISIBILITY

and CIVIC DESTABILIZATION

- Our Civilizational Challenge -

PART I:

SIX RINGS of DIGITIAL INVISIBILITY and CIVIC DESTABILIZATION

- OUR CIVILIZATIONAL CHALLENGE -

Preface: Past to Future: "Benjamin Franklin's Civic Logic and Collective Action: Timely Notice of Public Danger"

Chapter 1: Problem's Description and Model:
- Description: "Our Digital Invisibility Crisis: The Currents Pulling Humanity Under" (Creative Essay)
- Diagnostic Model: "Six Rings of Digital Invisibility and Civic Destabilization"

Chapter 2: Problem's Technical and Interpretive Contexts:
- Technical Context: "Seven Environmental Stressors of the Civic Digital Field"
- Interpretive Context: "The Public's Right to Parity of Civic Sight"

Chapter 3: Problem's Domestic and Foreign Societal Contexts:
- Domestic: "Our Governments' Digital Divide: State AI Initiatives Collide with a Federal Challenge"
- Foreign: "Our Foreign-Origin Civic Digital Crisis: Categorically Different Than Military-Monitored Cyber Warfare"

Chapter 4: Our Interim Mitigation Opportunities:
Three Practical, Cost-Effective Proposals as We Replace Agentic AI with Non-Agentic Civic AI in U.S. Government:
- Legislation to Limit Our Exposure: "Six Rings Zero Acceleration Requirement Proposal for All U.S. Agentic AI: A Civic Hippocratic Oath for U.S. Agentic AI ((A Civic AI Visibility Proposal for U.S. Government at All Levels)"
- Cross-State Coordination: "A Nation Under Civic Digital Stress and One Standardized Rider That Could Help Tie the States Together: Non-Agentic Cross-State Civic Digital Visibility Rider (Legislative Proposal Inspired by Benjamin Franklin's Creation of the American Postal System)"
- Demanding Non-Agentic Future Options for Procurement: "Shouldn't We Have Non-Agentic Civic AI Options for U.S. Government at All Levels? - Practical and Constitutional Perspectives for U.S. Citizens and Government Technology Procurement Professionals"

Afterword: What Long-Term Remedies Might a Democratic Society Allow?

Preface: Past to Future: Benjamin Franklin's Civic Logic and Collective Action

Timely Notice of Public Danger

"It is of great importance to the Public to give early Notice of Danger." - Benjamin Franklin, 1773

Danger rarely arrives as a trumpet blast. It appears first as a pattern—uneven, partial, easy to dismiss.

A rumor passed quietly between neighbors. A subtle shift in numbers on a ledger. A system beginning to behave strangely, but not yet catastrophically. Those who notice it earliest often hesitate, unsure whether what they see is real, whether it is their place to speak, or whether anyone else will see it too.

Conscience may stir at the sight of approaching harm, but conscience alone does not create safety. A warning unheard is indistinguishable from silence. A truth unshared does not yet belong to the public.

Franklin understood this. Long before sirens, sensors, and national alert systems, he treated early notice of danger as a civic obligation rather than a private anxiety. To see danger and remain quiet was not neutrality; it was a failure of coordination.

But notice is only the beginning.

Threats that exceed any one person - fires, epidemics, disinformation, systemic injustice - cannot be met by solitary civic/moral resolve. They require alignment: neighbors willing to recognize the same reality, trust the same signals, and act together before catastrophe forces the issue.

Franklin did not romanticize unity. He did not promise harmony or moral perfection. He offered something more practical and more demanding: the recognition that fragmentation has consequences, and that delay is itself a decision.

For this, he chose an image that did not flatter its audience, did not explain itself, and did not soften its meaning.

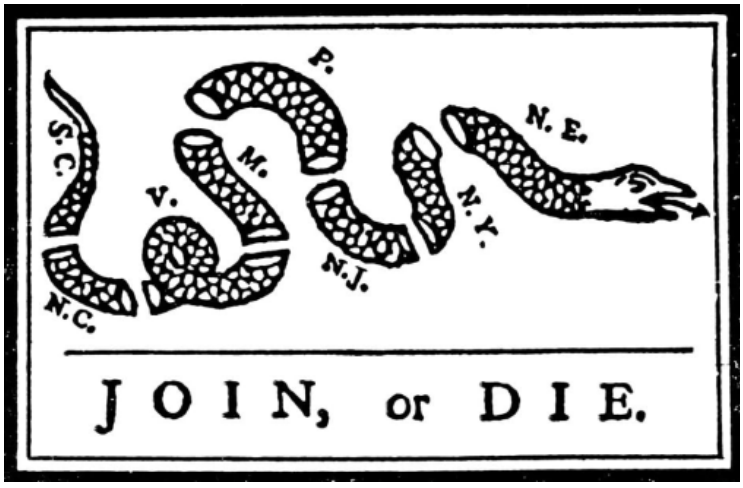

Benjamin Franklin's

"Pennsylvania Gazette,"

1754

Chapter 1: Problem's Description and Model

Description:

- "Our Digital Invisibility Crisis: The Currents Pulling Humanity Under" (Creative Essay)

Diagnostic Model:

- "Six Rings of Digital Invisibility and Civic Destabilization"

Our Digital Invisibility Crisis: The Currents Pulling Humanity Under (Creative Essay)

Today, a new digital river courses through human life - swift, deep, and indifferent. It carries our conflicts, myths, fears, and frailties faster and farther than any society has previously endured. This river does not create our ills, but it magnifies them, sweeping consequences across individuals, institutions, and nations before we can fully see where we are being carried.

Individually, we increasingly face the loss of visibility at the speed of modern life.

Collectively, our democracies increasingly face blindness, as speed, scale, and fragmentation overwhelm human perception. The government cannot protect what it cannot perceive.

And so we begin here - because no alignment is possible unless we first understand the currents already pulling us under.

The Currents Pulling Humanity Under (Creative Essay)

We face a crisis not of darkness, but of blindness - not the extinguishing of civic/moral light, but the accelerating loss of our ability to see clearly enough, quickly enough, to govern responsibly.

With the advent of AI, the world has not stopped. It has accelerated. And in that acceleration, something essential has slipped from our grasp: the visibility required for democratic self-governance.

Across schools, hospitals, courtrooms, neighborhoods, and national elections, the instruments of public life now operate at speeds and scales that exceed human perception.

We are attempting to navigate storms we cannot fully see, using maps drawn for calmer waters. What follows is not a prediction. It is an observation. It is documentation.

Unseen currents surround our public life - shifting their courses beneath our perception.
Accelerating forces that outpace our sight.

A school counselor sits alone after dismissal, staring at a list of students flagged by an automated system. Each name represents a child navigating the fragile terrain of adolescence. But the system offers no reasoning - only scores, only urgency.

She wants to act responsibly. She wants to serve each child with the care their situation demands, but the logic that generated these alerts remains sealed from her. This is not an isolated failure.

In America's 100,000 public schools, educators increasingly depend on automated tools whose reasoning they cannot audit. The burden grows heavier. The clarity grows thinner. Those who must carry the civic/moral weight - who look children in the eye and make decisions that shape their futures - find themselves governing in fog.

Individual civic/moral overload is the first ring of destabilization. It begins quietly - with one overwhelmed professional. But it does not remain individual for long.

Pulling us all - quietly at first,
with a gradual strength
we don't recognize until we are already drifting.

A pastor, a teacher, and a hardware store owner meet to discuss vandalism in their neighborhood. Once, they would have shared a common understanding of the facts. A conversation would have led toward unity.

But now, each arrives carrying a different version of the same events - curated by algorithmically tailored feeds optimized for engagement, not accuracy.

They speak in good faith. They care about their community. Yet they leave further apart than when they arrived. Trust does not collapse in an explosion. It exits the room quietly.

This is community fragmentation - the second ring.

Platforms prioritize reaction over reflection, engagement over truth. The informational commons dissolve into a thousand isolated realities. When communities cannot coordinate meaning, civic institutions lose their anchor.

Pulling us away - from conscience expression and values once shared.

A judge prepares to sentence a defendant. On her screen, a number appears: a risk score generated by an algorithm neither she nor the attorneys can fully explain. It offers authority without transparency, precision without reasoning.

She has spent decades learning to judge carefully. Now, for the first time, she feels her role shifting - from decision-maker to interpreter of invisible machinery.

This is institutional blinding - the third ring.

Hospitals, housing authorities, public benefits agencies - all increasingly operate through tools whose inner workings remain obscure even to those who use them.

A 2022 Government Accountability Office review found that more than half of federal agencies deploy artificial intelligence without adequate transparency. Institutions designed for daylight now operate in digital fog. And when governance systems cannot explain their own decisions, legitimacy erodes.

Pulling us apart – our communities, our civic institutions.
Unraveling the bonds that once held us in common purpose.

On the eve of a national election, a video surfaces that appears to show a leading candidate committing a crime. The lighting is convincing. The voice is authentic. Within ninety minutes, millions have seen it.

By the time analysts confirm it is a deepfake, the narrative has metastasized across the informational landscape. Verification requires hours. Viral spread requires minutes.

This is national polarization accelerated beyond democratic response - the fourth ring.

When falsehood outruns truth, public deliberation becomes impossible. When synthetic media becomes indistinguishable from an authentic recording, shared reality fractures. Democracy cannot deliberate if it cannot discern.

Pulling us downward - or quietly outward toward dangers unknown.
Both escaping awareness 'til pulled in harm's way.

A Chicago physician begins her shift. Patient records load slowly. A colleague mentions a ransomware attack in Singapore. Minutes later, her own systems begin to falter. She has not been attacked. She has been connected.

We now inhabit a world where digital infrastructure links institutions across continents, compressing global shocks into local failures.

This is the fifth ring: global risk compression.

Local vigilance cannot secure systems woven into a global digital fabric. A tear in one region becomes a crisis in another. In 2023 alone, healthcare ransomware disruptions cost American institutions an estimated $78 billion. No community can protect itself in isolation.

And the currents are growing stronger.

A fringe political group releases a provocative claim online. At first, no one notices. But an AI tool begins generating thousands of narrative variants, testing each for emotional impact, targeting audiences most susceptible to its message.

Within hours, the narrative appears organic. Within days, it influences public sentiment.
And the origin becomes untraceable.

This is the sixth ring: cascading tech-enabled amplifiers.

By 2025, researchers estimate that up to 30% of political content may be synthetic or manipulated—beyond the reach of reliable detection.

We have entered an environment where baseline reality itself wavers. This is not a condition that can be solved by trying harder. It is a structural shift that exceeds unaided human perception.

Unseen currents surround our public life - shifting their courses beneath our perception.
Accelerating forces that outpace our sight.
Pulling us all - quietly at first,
with a gradual strength
we don't recognize until we are already drifting.

Pulling us away - from conscience expression and values once shared.

Pulling us apart – our communities, our civic institutions.
Unraveling the bonds that once held us in common purpose.

Pulling us downward - or quietly outward toward dangers unknown.
Both escaping awareness 'til pulled in harm's way.

And the currents are growing stronger.

What you have just read is not a thought experiment. It is a record of experiences that millions of people are already living inside, often without shared language for what is happening to them. The school counselor, the pastor and neighbors, the judge, the physician, the voters scrolling through synthetic storms—each is encountering the same underlying currents before they are named. These are not isolated stories, but recurring patterns: overload, fragmentation, institutional blinding, amplification, governance strain, and cascading consequence.

In the next chapter, we'll step back from our digital river's surface and begin to consider these patterns further - to make them visible.

Once we can visualize the often invisible, but always consequential, conditions shaping our civic digital environment, we can begin to consider how a democratic people might respond.

Six Rings of Digital Invisibility and Civic Destabilization

Including

Franklin-Inspired Analysis Between Diagnosis and Design

(Public Utility and Cost Transfer Considerations)

Across politics, public health, education, and everyday civic life, many people sense that something fundamental is off—but struggle to name it clearly. Institutions appear reactive. Public trust is strained, and crises feel more frequent, more confusing, and harder to govern than they once did.

This essay offers a simple framework to help us understand why: the Six Rings of Digital Invisibility and Civic Destabilization.

The Six Rings are not an attack manual. They are civic-scale visibility failure conditions—the progressive breakdown of democratic perception and coordination under machine-speed foreign-origin digital threats.

Rather than treating today's civic digital disruptions as isolated failures, the Six Rings describe a layered and cascading civic condition: a progressive breakdown in our shared ability to see, interpret, and coordinate responsibly in a world shaped by digital scale, speed, and complexity.

The problem is not a loss of intelligence or goodwill. It is a growing mismatch between civic/moral responsibility and perceptual capacity.

Moving from individual overload toward systemic cascade, the Six Rings trace how strain accumulates—first in human attention, then in visibility, institutions, governance, and finally across interconnected systems. Each Ring builds upon the conditions established before it, revealing how civic destabilization can emerge even when people and institutions are acting in good faith.

What We Mean by Digital Invisibility

Digital invisibility does not mean that harm itself cannot be seen. In many cases, harm is visible in its immediate effects. What becomes difficult to see is the shared structure behind those effects: how harm arises, who is affected, how consequences accumulate, and where meaningful intervention remains possible.

Across the Six Rings, harms may accumulate in plain sight while remaining practically invisible to those responsible for governing, repairing, or responding to them.

This invisibility is not necessarily the result of secrecy or malicious intent. It can emerge from the way data, incentives, interfaces, and institutional responsibilities are structured—fragmenting causes and consequences across time, systems, and jurisdictions.

The harm may therefore remain visible while the pattern producing it remains invisible.

How the Six Rings Progress

The Six Rings describe a recurring progression in which digitally mediated systems make consequential civic risks progressively harder to see, judge, and correct.

The Rings are sequential but overlapping. Later Rings do not replace earlier ones; they compound them.

The progression begins with civic/moral overload, as individuals and institutions encounter more signals, demands, and competing claims than human attention can meaningfully process. It moves toward civic invisibility, as important harms and causal relationships fall outside shared public visibility. Institutional blindness and normalization follow when fragmented signals become accepted as sufficient for continued operation. Amplification without judgment then accelerates effects according to digital performance measures rather than civic evaluation. Governance under strain emerges as established institutions struggle to respond to fast-moving, distributed, and cumulative conditions.

Finally, these accumulated conditions can produce civic destabilization in plain sight: trust erodes, coordination falters, and democratic institutions continue operating under increasingly diminished shared visibility.

The Rings therefore describe not six isolated problems, but a compounding civic process.

RING I – CIVIC/MORAL OVERLOAD

When Conscience Faces More Than It Can Carry

The first destabilizing pressure is not a failure of civic/moral intent. It is a failure of civic/moral capacity under load.

Modern citizens are asked, daily, to absorb, interpret, and emotionally process volumes of human suffering, ethical controversy, and existential risk that no previous generation was built to carry alone. War, climate disruption, political fracture, public health crises, institutional failure, economic precarity, and algorithmically amplified outrage now arrive continuously, compressed into the same cognitive and emotional space once reserved for family, work, and local community. This does not make people callous. It makes them overwhelmed.

Human conscience evolved to respond to situated civic/moral demands - those that could be seen, named, and acted upon within a limited horizon. Today, that horizon has exploded: civic/moral claims arrive faster than reflection, wider than response, and without clear pathways for action. The result is not apathy, but civic/moral saturation: a state in which the internal signal of "ought" is constantly triggered without resolution.

Over time, predictable consequences follow:

- Ethical fatigue replaces deliberation
- Urgency crowds out reflection
- Emotional numbing substitutes for judgment
- Civic/Moral certainty hardens where nuance once lived

AI-generated deepfake videos and images flooded 2025 election campaigns worldwide, complicating efforts to verify authentic political communication (Firstpost, 12/2025).

Officials in over 50 countries reported AI-generated content as a factor in recent elections, from Canada to Germany and Portugal (New York Times, 06/2025).

Election experts warned that AI-generated audio can now spoof official voter hotlines, confusing citizens about polling locations and rules (Brennan Center, 05/2024).

Media analysts noted a steady flow of low-quality AI-generated content entering social feeds, blurring boundaries between authentic reporting and synthetic material (Poynter, 12/2025).

Field experiments showed AI-generated misinformation reduces overall trust in news, even among audiences relying on reputable outlets (Digital Content Next, 09/2025).

Reuters Institute surveys found widespread uncertainty about when AI systems are being used in news and digital platforms (Reuters Institute, 10/2025).

Researchers observed that awareness of AI-generated misinformation leads some users to disengage from news altogether (Reuters Institute, 10/2025).

People begin to oscillate between anxiety and withdrawal, outrage and exhaustion, conviction and paralysis. None of these states reflect a loss of conscience. They reflect a system demanding more civic/moral processing than the human nervous system can sustain.

This overload is intensified by digital environments that do not distinguish civic/moral relevance from engagement value. Tragedy, injustice, and fear are delivered alongside entertainment, opinion, and speculation, flattened into a single attention stream, and the mind is given no time to metabolize one civic/moral claim before the next arrives.

In this condition, conscience does not disappear. It short-circuits.

Well-intentioned individuals may retreat into simplifications - not because they reject civic/moral responsibility, but because complexity becomes unbearable. Others disengage entirely, not out of indifference, but as an act of self-preservation, while still others cling to rigid frameworks that promise certainty, even at the cost of accuracy or compassion.

These are not personal failures. They are systemic stress responses. Ring I matters because every subsequent destabilization builds upon it: a civically/morally overloaded population struggles to notice subtle harms, to sustain long-term concern, or to coordinate responsibly with others. When conscience is constantly triggered but rarely resolved, it loses its capacity to guide.

Analysts warned that low-cost generative AI tools now enable small groups to mount influence operations previously requiring state-level resources (Brennan Center, 05/2024).

International observers tracked AI-assisted interference and influence operations becoming a routine feature of contemporary elections, not a one-off anomaly (CIGI, 08/2025).

Political strategists increasingly relied on AI content generation for attack ads and rapid-response messaging in 2025 national races (Economist, 07/2025).

Media observers reported no single “deepfake apocalypse,” but a steady accumulation of AI-generated content degrading information quality over time (Poynter, 12/2025).

Survey research showed many citizens worry AI will worsen misinformation and increase confusion about what to trust (Reuters Institute, 10/2025).

Researchers found that realistic AI-synthesized images significantly increase belief in false headlines (Harvard Misinformation Review, 11/2025).

Experimental studies suggested corrections have limited effect when false stories are paired with highly realistic AI-generated images (Harvard Misinformation Review, 11/2025).

This is the first fracture in civic visibility: not the dimming of civic/moral light, but the exhaustion of the bearer.

Before misinformation spreads, before institutions fail, before amplification distorts reality, conscience itself is placed under unsustainable strain. And when the load exceeds the carrier, even the strongest flame begins to flicker - not from weakness, but from weight.

RING I – CIVIC CAPACITY OVERLOAD

America's Current Related Costs

When individuals are routinely required to absorb volumes of risk, suffering, and ethical complexity that exceed ordinary human capacity, the resulting strain does not remain private. Civic capacity overload contributes to rising healthcare expenditures, workplace error and absenteeism, declining institutional trust, preventable public-health failures, and increased reliance on emergency services and crisis interventions. These costs are already borne collectively through public systems, employers, insurers, and local governments — often without being recognized as the downstream effects of sustained cognitive and civic overload. What appears as individual burnout or disengagement frequently reappears later as shared fiscal and institutional burden, measured in billions over time.

Governments explored AI tools to analyze public comments at scale for regulatory consultations, reshaping civic participation workflows (OECD, 09/2025).

An OECD report documented governments deploying AI chatbots for citizen services, raising questions of accountability and error correction (OECD, 09/2025).

US guidance urged agencies to conduct impact assessments before deploying high-risk AI systems affecting rights or access to services (GDPR Local, 09/2025).

Multiple US states created AI task forces to study risks to civil rights, workforce impacts, and public safety before large-scale deployment (GovTech, 11/2025).

State legislatures emphasized transparency and human oversight as cornerstones of responsible AI use in public services (GovTech, 11/2025).

Congressional debates highlighted tensions between national AI standards and states' desire to act as laboratories of democracy (StateScoop, 12/2025).

Commentators warned that public trust in AI governance may depend as much on visibility and accountability as on technical performance (World Economic Forum, 07/2025).

Franklin-Inspired Analysis

(Public Utility and Cost Transfer Considerations)

From a public-works perspective, recurring overload signals a design failure rather than a personal one. When shared systems consistently impose demands that exceed normal human processing limits, costs are quietly transferred until they surface elsewhere as public expense. A civic response does not attempt to reform individual character or automate judgment; instead, it seeks to reduce unnecessary load by restoring visibility, pacing, and proportionality through shared, non-coercive infrastructure. By shifting a portion of this burden away from isolated individuals and into neutral civic tools that clarify conditions without directing decisions, societies can lower predictable downstream costs while preserving personal agency. As with earlier public works, the justification rests not on civic/moral improvement, but on reducing recurring civic strain and stabilizing essential social function.

RING II – CIVIC/MORAL INVISIBILITY

(When Harm Occurs Beyond the Reach of Witness)

If Ring I strains conscience through excess, Ring II undermines it through absence.

Civic/Moral judgment depends not only on intention, but on visibility - the ability to perceive who is affected, how harm unfolds, and where responsibility lies. When consequences cannot be seen, conscience cannot reliably orient action, no matter how sincere the civic/moral impulse.

US and European regulators debated how AI use in derivatives markets could introduce "too fast to stop" systemic financial risks (CRS, 12/2024).

A GAO-cited study warned that concentrated reliance on a small number of AI providers could magnify instability in financial markets (CRS, 12/2024).

Researchers highlighted "financial deepfakes," where AI-generated media can move markets by spreading false corporate or macroeconomic news (Regulatory Review, 11/2025).

A widely shared AI-generated image falsely depicting a Pentagon explosion briefly erased hundreds of billions in market value before being debunked (Regulatory Review, 11/2025).

Central bankers and IMF officials flagged AI-driven trading as a contributor to faster and harder-to-explain volatility episodes (Galapagos Capital/IMF, 04/2025).

Analysts estimated that AI-linked trading accounts now handle a majority of US equity transactions, complicating market oversight (Regulatory Review, 11/2025).

Experts warned that AI-based trading bots could inadvertently learn collusive strategies resembling price-fixing or bid-rigging (Liquidity Finder, 02/2025).

For much of human history, moral harm was largely proximate. Suffering had faces. Decisions carried observable effects, and responsibility could often be traced through direct human relationships or local institutions. While injustice certainly existed, its pathways were more legible.

That condition no longer holds.

Today, harm increasingly occurs at a distance - mediated by platforms, obscured by scale, fragmented across jurisdictions, and delayed in time. Digital systems allow actions taken in seconds to reshape lives elsewhere without the actor ever encountering the affected person - or even realizing an effect occurred.

This produces civic/moral invisibility: situations in which real harm exists, but no clear civic/moral signal reaches those positioned to intervene.

Civic/Moral invisibility commonly appears in four forms:

- Distance without encounter: People are affected by decisions made far away, by individuals or systems they will never meet.
- Scale without story: Suffering appears statistically rather than personally, flattening human impact into abstraction.
- Anonymity without accountability: Actions taken behind screens dissolve the felt presence of others and weaken restraint.
- Delay without feedback: Consequences surface long after decisions are made, severing cause from civic/moral recognition.

NPR reported growing concern about AI-amplified misinformation campaigns targeting retail investors with fake earnings or merger stories (NPR, 10/2025).

Reports documented a sharp rise in deepfake-enabled fraud schemes affecting corporate finance and investor confidence (Regulatory Review, 11/2025).

The IMF highlighted risks as financial institutions outsource core analytical functions to a small number of AI vendors (Galapagos Capital/IMF, 04/2025).

Commentators described emerging AI-related market threats as increasingly opaque to traditional risk management tools (Regulatory Review, 11/2025).

A congressional brief noted that AI-driven models could accelerate contagion in derivatives markets during stress events (CRS, 12/2024).

Analysts reported a sharp increase in AI-related capital markets patents, signaling rapid automation across financial services (Galapagos Capital/IMF, 04/2025).

Corporate security experts warned that AI-enabled impersonation of executives is increasingly used to authorize fraudulent high-value transfers (World Economic Forum, 07/2025).

In such conditions, conscience is not overridden — it is starved of information.

People may sincerely wish to do no harm while unknowingly participating in systems that cause it. Others may witness fragments of harm but lack the context needed to understand its source or magnitude, and institutions may act with partial sight, addressing visible symptoms while deeper causes remain hidden.

This invisibility erodes civic/moral coordination: when individuals cannot see how their actions connect to broader outcomes, shared responsibility fractures. Each person experiences only a sliver of the civic/moral landscape, while the whole remains obscured.

Over time, a dangerous distortion sets in: harm begins to feel accidental, diffuse, or inevitable — rather than authored, traceable, and contestable.

Importantly, civic/moral invisibility does not require malice, deception, or conspiracy. It emerges naturally when systems grow faster, larger, and more interconnected than human perception can track. Digital mediation intensifies this effect by stripping away the sensory and relational cues that once anchored civic/moral awareness.

A person may believe they are acting responsibly while the consequences of their actions accumulate elsewhere — unseen, unacknowledged, and unaddressed.

An MIT study estimated that AI could already perform tasks equivalent to 11.7% of US jobs, particularly in finance and healthcare (CNBC, 11/2025).

Yale researchers found no broad labor-market collapse but noted localized disruption in occupations highly exposed to AI tools (Yale Budget Lab, 08/2025).

World Economic Forum estimates suggested roughly 60% of jobs in advanced economies may be affected by AI-driven task reallocation (J.P. Morgan, 06/2025).

Analysts described a “white-collar recession,” with firms citing AI efficiencies while reducing professional and technical roles (Cresset Capital, 11/2025).

Gartner data indicated fewer than 1% of global layoffs in 2025 were directly attributable to AI productivity gains (Cresset Capital, 11/2025).

Job-cut tracking showed nearly 950,000 US layoffs through September 2025, with AI frequently referenced in corporate restructuring narratives (Cresset Capital, 11/2025).

Labor economists emphasized that AI’s redistribution of tasks within jobs may prove more disruptive than outright job elimination (J.P. Morgan, 06/2025).

Ring II deepens the fracture introduced in Ring I. A civically/ morally overloaded population now faces a world where even sustained attention cannot reliably reveal where civic/moral effort is needed most, and conscience strains not only under weight, but under blindness.

When harm becomes invisible, responsibility weakens - not because people no longer care, but because caring loses its bearings.

This is the second collapse of civic visibility: a world in which civic/moral signals fade before they can guide judgment, leaving conscience sincere but disoriented.

RING II - CIVIC INVISIBILITY

America's Current Related Costs

When harm occurs beyond the reach of ordinary perception, corrective action is delayed or misdirected. Civic invisibility contributes to preventable public expenditures across healthcare, environmental response, consumer protection, labor regulation, and digital enforcement, as harms accumulate unnoticed until they require costly intervention.

Distance, scale, anonymity, and delayed feedback allow negative consequences to spread without early correction, transforming manageable issues into systemic liabilities. These costs are absorbed collectively - often in the form of remediation, litigation, regulatory expansion, and crisis response - representing substantial public expense that could have been reduced through earlier visibility.

Reuters Institute data indicated that a growing share of younger audiences use AI assistants to access news, altering how information is filtered and summarized (Reuters Institute, 10/2025).

An international audit of AI assistants found recurring hallucinations and misrepresentation when summarizing news articles (EBU/Reuters Institute, 10/2025).

Survey research showed that uncertainty about AI use in media contributes to declining trust in news across multiple countries (Reuters Institute, 10/2025).

Media analysts reported that AI-generated content increasingly blurs distinctions between reporting, aggregation, and synthetic material (Poynter, 12/2025).

Publishers experimented with AI transcription and summarization tools while retaining human editors to preserve credibility and accountability (Poynter, 12/2025).

Several news organizations signed licensing agreements with AI firms to control training data and stabilize revenue models (Poynter, 12/2025).

Researchers observed that exposure to AI misinformation can reduce trust even in reputable outlets not directly involved in false reporting (Digital Content Next, 09/2025).

Franklin-Inspired Analysis

(Public Utility and Cost Transfer Considerations)

From a public-works standpoint, invisibility is a structural problem, not a civic/moral one. When systems operate beyond the range of human sight, responsibility weakens and costs compound. Civic/moral infrastructure has historically addressed such conditions by extending shared visibility - making risks legible before they metastasize. A civic response does not assign blame or automate judgment; it restores informational continuity between action and consequence. By illuminating patterns that are otherwise dispersed across time, distance, and institutions, shared tools can reduce delayed harm and lower the long-term public costs associated with reactive correction.

RING III – INSTITUTIONAL INVISIBILITY

When Systems Act Beyond Human Oversight

As civic/moral invisibility spreads at the individual level, a parallel failure emerges within the very institutions meant to coordinate responsibility. Ring III marks the point at which institutions themselves lose visibility into the systems they operate.

Modern institutions — governments, hospitals, corporations, regulatory bodies, and public services — were designed for worlds in which causal chains were slower, narrower, and more legible. Authority rested on the assumption that decision-makers could reasonably understand the systems under their control.

Security analysts documented attackers using AI to generate polymorphic malware that evades traditional signature-based defenses (SecurityBrief, 12/2025).

Reports described AI-assisted cyberattacks crafting dynamic command-and-control traffic that blends into legitimate network activity (SecurityBrief, 12/2025).

Cybersecurity experts coined the term “living off the LLM” to describe attackers exploiting AI-native system features during intrusions (SecurityBrief, 12/2025).

Industry analysts highlighted AI-enabled cyberattacks targeting water facilities, airports, and automotive systems during 2025 (SecurityBoulevard, 12/2025).

CISA guidance urged operators to segregate AI systems from operational technology networks in critical infrastructure environments (SecurityBrief, 12/2025).

DeepStrike data linked increased infrastructure exposure to rapid integration of cloud-based AI services without strong segmentation controls (DeepStrike, 10/2025).

Regulators considered mandatory reporting of AI-related cyber incidents to improve situational awareness across critical sectors (SecurityBoulevard, 12/2025).

That assumption is now under strain.

Contemporary institutions increasingly rely on complex, interdependent systems - digital platforms, algorithmic processes, global supply chains, and automated decision tools - whose behavior no single actor fully understands. Authority remains centralized, but comprehension has become distributed, fragmented, or absent altogether.

This produces institutional invisibility: conditions in which organizations exercise power without clear sight of how outcomes are generated, where risks accumulate, or when harm is emerging.

Institutional invisibility commonly appears as:

- Opaque system behavior: Decisions arise from interactions between technical systems, policies, and incentives that cannot be easily traced or explained.
- Fragmented responsibility: No single role or office possesses full situational awareness, even as outcomes affect many.
- Lagging oversight: Harms surface only after they have compounded, when intervention is costly or politically fraught.

Analysts observed that AI-assisted cyber defense tools are reshaping how security teams prioritize alerts and respond to threats (SecurityBoulevard, 12/2025).

Studies documented AI tools identifying software vulnerabilities earlier in development, potentially reducing some categories of cyber risk (IndustrialCyber, 11/2025).

Policy briefs warned that outages driven by AI-enabled cyber incidents could undermine public confidence in digital utilities (DeepStrike, 10/2025).

Government guidance flagged AI deployment in critical infrastructure as requiring specialized safety and resilience frameworks (GDPR Local, 09/2025).

Consumer advocates cautioned that opaque AI systems in infrastructure management complicate accountability during service failures (DeepStrike, 10/2025).

Analysts noted that as AI mediates core services, small errors can propagate rapidly across interconnected systems (OECD, 09/2025).

International organizations emphasized that maintaining trust in digital infrastructure depends on visibility into AI decision-making processes (OECD, 09/2025).

- Procedural substitution for judgment: Rules and models stand in for understanding when systems become too complex to interpret directly.

In such environments, institutions continue to function - but increasingly on trust in process rather than understanding of consequences.

This is not institutional negligence. It is institutional overload.

As systems accelerate, organizations respond by layering procedures, metrics, and compliance structures atop one another. These measures create the appearance of control while often further obscuring causal pathways. Over time, leaders may sense that something is wrong without being able to say precisely where or why.

The civic/moral danger here is subtle but profound: when institutions cannot see clearly, they tend to prioritize defensibility over truth, continuity over correction, and stability over accountability. Harm is treated as a systems issue rather than a civic/moral signal, and responsibility diffuses into committees, dashboards, and protocols.

Citizens sense this disorientation. Public trust erodes not only because institutions err, but because they cannot convincingly explain their actions. Transparency becomes performative, offering outputs without insight, and accountability becomes retrospective, focused on blame after the fact rather than prevention before harm

Hospitals increasingly adopted AI tools for patient monitoring, imaging analysis, and clinical risk prediction during 2025 (Patient Safety Journal, 10/2025).

Event analyses found AI-assisted imaging systems improved detection accuracy in some cases but did not consistently prevent downstream patient harm (Patient Safety Journal, 10/2025).

ECRI named AI-enabled clinical decision-support systems the top health technology hazard for 2025 (ECRI, 12/2024).

Patient safety experts warned that poorly overseen AI tools can introduce new diagnostic and treatment risks alongside efficiency gains (Health Journalism, 12/2024).

US policymakers tracked more than 250 state-level bills regulating AI use in health care, addressing safety, privacy, and bias concerns (Manatt Health, 12/2025).

The FDA issued draft guidance signaling closer oversight of adaptive AI algorithms embedded in medical devices (S3 Connected Health, 12/2025).

A UK proposal outlined an AI early-warning system to flag emerging patient safety risks across the NHS using near–real-time data (Patient Safety Learning, 07/2025).

Ring III intensifies the earlier fractures. Individuals already strained by civic/moral overload and civic/moral invisibility now encounter institutions that cannot reliably illuminate the systems shaping their lives, and the gap between authority and understanding widens, creating fertile ground for cynicism, polarization, and withdrawal.

At this stage, the failure is no longer simply one of perception. It is a failure of institutional sight.

Institutions continue to decide, allocate, and enforce - but increasingly without the shared visibility required to justify those actions to the public or to themselves.

This is the third collapse of civic visibility: when the structures designed to coordinate civic/moral responsibility lose their capacity to see the consequences of their own operation.

RING III – INSTITUTIONAL BLINDNESS

America's Current Related Costs

As institutions rely on increasingly complex and interdependent systems, failures of institutional visibility produce measurable downstream costs. Policy errors, misallocated resources, delayed responses, and ineffective oversight generate waste, public mistrust, and preventable harm that must be later addressed through audits,

Hospitals reported using AI tools to reduce medication errors and sepsis-related harms, integrating alerts into clinical safety dashboards (American Hospital Association, 12/2025).

Researchers found clinicians frequently override AI recommendations, with trust levels strongly influencing adoption outcomes (World Economic Forum, 12/2025).

Analysts emphasized that perceived safety and transparency of AI tools shapes health system uptake as much as technical performance (World Economic Forum, 12/2025).

Public health researchers cautioned that AI-generated visual misinformation may exacerbate vaccination hesitancy and health behavior challenges (Harvard Misinformation Review, 11/2025).

Patient advocacy groups urged transparent reporting of AI-related clinical incidents to prevent repeat errors across health systems (Parker Poe, 11/2025).

State regulators embedded civil rights protections into health AI rules, including audit requirements and patient appeal mechanisms (Manatt Health, 12/2025).

Health data experts flagged privacy and governance risks tied to expanding AI use in electronic health records (NCBI Bookshelf, 03/2025).

investigations, litigation, and emergency interventions – often at significantly higher expense than earlier corrective action would have required.

When decision-making authority outpaces institutional comprehension, errors become expensive - not only financially, but in lost legitimacy and coordination capacity. These costs are distributed across taxpayers, service users, and frontline institutions, often surfacing only after corrective options have narrowed.

Franklin-Inspired Analysis

(Public Utility and Cost Transfer Considerations)

Franklinian reasoning treats institutional blindness as a signal that systems have outgrown their interpretive supports. Civic infrastructure historically emerged to help institutions see what scale and speed made opaque, not to replace judgment but to inform it. A civic response focuses on restoring institutional sight through shared, intelligible visibility rather than adding procedural complexity. By reducing reliance on after-the-fact correction and defensive governance, such infrastructure lowers predictable public expense while strengthening institutional function and public trust.

RING IV — AMPLIFICATION WITHOUT JUDGMENT

(When Scale Outruns Discernment)

If the earlier rings strain conscience, obscure harm, and blind institutions, Ring IV introduces a more volatile force: amplification divorced from judgment.

The Brennan Center warned that predictive policing tools can magnify existing bias and chill constitutionally protected activity (Brennan Center, 11/2025).

Analysts reported limited independent evidence that predictive policing systems significantly improve public safety outcomes (Brennan Center, 11/2025).

Legal commentary cautioned that AI surveillance tools may struggle to distinguish victims from suspects, raising due process concerns (Jones Walker, 10/2025).

Researchers suggested AI-based crime prediction would require dramatically higher accuracy thresholds to justify operational reliance (Jones Walker, 10/2025).

Civil liberties advocates called for public inventories of AI systems used in criminal investigations to support transparency and accountability (Jones Walker, 10/2025).

Policing technology firms acknowledged risks of algorithmic bias when models are trained on historically skewed enforcement data (SoundThinking, 05/2025).

Experts recommended continuous bias audits and mandatory human verification before acting on AI-generated policing leads (SoundThinking, 05/2025).

Digital systems excel at scale. They replicate, distribute, and intensify signals with extraordinary speed and reach. What they do not do, by default, is distinguish wisdom from distortion, care from coercion, or truth from emotional potency. Signals that travel well are amplified; signals that require patience, context, or proportion are not.

This is not a design failure so much as a structural reality. Systems optimized for engagement, efficiency, or reach amplify whatever travels well through networks, not whatever deserves deliberation.

As a result, civic/moral signals are no longer merely communicated. They are accelerated.

In this environment, emotionally charged content outpaces reflective content, simplified narratives spread faster than careful ones, outrage outruns explanation, and fear travels farther than context. The volume of amplification is mistaken for importance, and repetition substitutes for verification.

Judgment cannot keep pace.

Amplification without judgment manifests in several destabilizing patterns:

- **Velocity over verification**: Claims spread faster than they can be evaluated or corrected.

Legal analysts described 2025 as the year hypothetical AI risks became operational legal questions, as courts confronted concrete failures and harms (Legal Commentary, 12/2025).

Scholars argued that existing constitutional doctrines are being tested by AI systems that mediate speech, surveillance, and access to public services (Harvard Law Review, 04/2025).

Courts began grappling with evidentiary standards for AI-generated material, including authenticity, provenance, and admissibility (Harvard Law Review, 04/2025).

Commentators warned that rapid AI adoption in administrative agencies risks outpacing procedural safeguards built into administrative law (Harvard Law Review, 04/2025).

Legal observers noted growing tension between efficiency gains from AI tools and due process protections in adjudicatory settings (Harvard Law Review, 04/2025).

Analysts emphasized that opacity in AI decision-making complicates judicial review and remedies when rights are affected (Harvard Law Review, 04/2025).

Policy debates highlighted uncertainty about liability allocation when automated systems contribute to harmful outcomes (Harvard Law Review, 04/2025).

- **Emotion over meaning** : Content optimized for reaction eclipses content designed for understanding.
- **Visibility over validity**: What appears most often is assumed to matter most.
- **Polarization through reinforcement**: Repeated exposure hardens positions before reflection can occur.

This dynamic places human conscience in a reactive posture. Instead of orienting action through deliberation, individuals and institutions respond to cascades already in motion, and civic/moral attention becomes event-driven rather than judgment-driven.

Importantly, amplification does not require falsehood to be destabilizing. Even accurate information, when stripped of proportion, context, or temporal pacing, can distort civic/moral perception. Tragedy without framing overwhelms and partial truths without synthesis mislead. Constant exposure without resolution exhausts.

Institutions struggle here as well. By the time a response is formulated, narratives have already solidified. Clarifications arrive too late and corrections feel defensive. Silence is interpreted as complicity, while restraint is mistaken for indifference.

Civil rights groups urged stronger disclosure requirements for AI use in government decisions affecting benefits, immigration, and eligibility determinations (OECD, 09/2025).

Public interest advocates warned that automated decision systems can entrench bias if audit and appeal mechanisms are weak or inaccessible (OECD, 09/2025).

An OECD analysis noted that AI-mediated service delivery may deter civic participation when citizens cannot understand or challenge outcomes (OECD, 09/2025).

Surveys showed persistent public unease about government use of AI in surveillance and law enforcement contexts (Reuters Institute, 10/2025).

Researchers found that fears of AI bias and opacity can reduce willingness to engage with digital public services (OECD, 09/2025).

Civil liberties organizations stressed the importance of notice and explanation when AI tools inform government action (Jones Walker, 10/2025).

Analysts warned that unequal access to appeals and explanations could widen existing civic participation gaps (OECD, 09/2025).

This creates a feedback loop. Amplification pressures institutions to react faster, which further degrades judgment and increases error, and each misstep fuels the next surge of attention.

Ring IV marks a critical acceleration point. The earlier collapses of visibility are no longer static conditions - they become dynamic instabilities. Distortion spreads not because no one cares, but because caring is continually hijacked by scale.

At this stage, civic/moral life becomes performative. Attention is demanded immediately, nuance is penalized, and reflection is crowded out by urgency. Conscience is no longer merely strained or blinded — it is pushed into motion before it can aim.

This is the fourth collapse of civic visibility: when the mechanisms that broadcast reality overwhelm the human capacity to judge it.

Franklin-Inspired Analysis

(Public Utility and Cost Transfer Considerations)

Unmoderated amplification accelerates distortion, conflict, and error faster than human judgment can respond. The fiscal consequences include crisis-driven policymaking, public-health misinformation costs, market instability, reputational damage requiring expensive counter-messaging, and repeated emergency mobilizations that could have been mitigated through earlier, slower and better-informed intervention.

Several US states advanced legislation requiring disclosure labels on AI-generated political advertisements and campaign materials (NCSL, 04/2025).

The “Protect Elections from Deceptive AI Act” was introduced to prohibit knowingly distributing materially deceptive AI media about federal candidates (Regula Forensics, 08/2025).

Lawmakers debated how to balance free expression protections with restrictions on deceptive synthetic political media (Regula Forensics, 08/2025).

State election officials reported challenges enforcing AI disclosure rules across rapidly evolving digital platforms (NCSL, 12/2025).

Congressional hearings examined whether centralized AI standards might preempt stronger state-level civil rights protections (StateScoop, 12/2025).

Policy analysts argued that preserving state experimentation could surface AI harms earlier than national-scale regulation alone (StateScoop, 12/2025).

Observers noted that public trust in democratic institutions increasingly hinges on transparency about how AI is used in governance (World Economic Forum, 07/2025).

Even accurate information, when amplified without proportion or context, can overwhelm institutions and citizens alike, increasing volatility and reactive spending. These costs are already borne across public and private systems, often invisibly, as resources are diverted from long-term stewardship to short-term containment.

From a civic perspective, amplification is not itself the problem; amplification without pacing or contextual visibility is. Public works have historically intervened where velocity outstripped human discernment, introducing stabilizing structures rather than suppressing activity. A civic response does not arbitrate truth or direct outcomes. It restores proportionality and temporal coherence, allowing judgment to re-enter the system before consequences cascade.

By reducing volatility-driven reaction, societies can lower recurring emergency expenditures and stabilize civic function without constraining expression or authority.

RING V – GOVERNANCE UNDER STRAIN

(When Authority Responds Without Shared Understanding)

By the time the first four rings are active, governance no longer operates in a stable informational environment. Ring V marks the point at which authorities must make consequential decisions among overload, invisibility, institutional blindness, and unchecked amplification, with no reliable common picture of reality to anchor their choices.

Intelligence analysts warned that AI-assisted influence operations are becoming a persistent feature of geopolitical competition rather than episodic campaigns (CIGI, 08/2025).

Security researchers documented state and non-state actors using generative AI to scale disinformation across multiple languages simultaneously (CIGI, 08/2025).

Defense analysts cautioned that AI-generated media complicates attribution and escalation control during international crises (World Economic Forum, 07/2025).

NATO-linked studies examined how AI-enabled information warfare blurs distinctions between civilian and military domains (World Economic Forum, 07/2025).

Analysts warned that low-cost AI tools allow smaller states and proxy actors to project disproportionate informational influence (Brennan Center, 05/2024).

Researchers noted that AI-assisted propaganda can be rapidly localized to exploit cultural and political fault lines (CIGI, 08/2025).

Policy briefs highlighted growing concern that AI-generated false signals could destabilize crisis communication channels (World Economic Forum, 07/2025).

Governance depends on more than authority. It depends on shared understanding - a minimum level of agreement about what is happening, what matters, and why certain actions are justified. When that shared ground dissolves, governance does not merely become harder, it becomes unstable.

In this condition, leaders face a narrowing set of options. Decisions must be made quickly, often with incomplete or conflicting information, under conditions of public scrutiny intensified by amplification. Delay is interpreted as weakness and caution is reframed as evasion. Every action, and every inaction, becomes a signal.

This produces governance under strain: a state in which institutions retain formal power but lose the informational and civic/moral conditions that make power legitimate and effective.

Common symptoms include:

- Reactive policymaking: Decisions are driven by immediate pressures rather than long-term coherence.
- Rule substitution for judgment: Fixed constraints replace deliberation when trust in interpretation collapses.
- Centralization of control: Authority tightens as a substitute for shared visibility.

International observers tracked AI use in cyber espionage targeting telecommunications and satellite infrastructure (IndustrialCyber, 11/2025).

Technology firms disclosed detecting attempts to adapt commercial AI models for intelligence collection and reconnaissance purposes (IndustrialCyber, 11/2025).

Analysts reported that AI-enabled cyber operations increasingly target supply chains rather than isolated systems (SecurityBrief, 12/2025).

Governments warned that AI-assisted attacks could compress decision timelines during cyber incidents affecting critical services (SecurityBrief, 12/2025).

Defense policy analysts cautioned that AI-mediated command-and-control systems raise risks of automation bias in security decision-making (World Economic Forum, 07/2025).

Studies suggested that reliance on opaque AI systems complicates civilian oversight of military and intelligence operations (World Economic Forum, 07/2025).

Experts emphasized the need for human-in-the-loop safeguards as AI tools enter national security workflows (World Economic Forum, 07/2025).

- **Legitimacy erosion:** Even well-intended actions are met with suspicion when rationales are unclear or contested.

Under such strain, governance oscillates between overreach and paralysis. Leaders may impose blunt controls to reassert order, or hesitate to act at all, fearing backlash from fragmented publics who no longer share a common picture of reality.

Importantly, this erosion is not confined to any single political system. Democratic governments struggle to sustain consent when citizens no longer agree on basic facts., while authoritarian systems tighten control when informational flows threaten stability. In both cases, governance reacts to uncertainty rather than resolving it.

As visibility collapses, governance begins to substitute control for comprehension.

This substitution carries civic/moral cost. When authority acts without shared understanding, citizens experience decisions as imposed rather than justified. Compliance weakens, trust decays, and the space for cooperative problem-solving shrinks, replaced by adversarial dynamics and symbolic politics.

International organizations debated norms for responsible AI use in defense, citing risks of arms-race dynamics (OECD, 09/2025).

UN working groups examined how AI may undermine existing confidence-building measures between rival states (UN, 12/2025).

Analysts warned that uneven AI governance standards across countries could create regulatory arbitrage and security gaps (OECD, 09/2025).

Diplomats noted growing tension between national AI sovereignty claims and the need for shared international rules (OECD, 09/2025).

Global surveys showed public concern that AI will concentrate geopolitical power in a small number of states and corporations (Reuters Institute, 10/2025).

Policy analysts argued that transparency and verification mechanisms are essential to sustaining trust in AI-related international agreements (World Economic Forum, 07/2025).

Observers cautioned that without shared visibility, AI risks could escalate silently across borders before diplomatic response is possible (OECD, 09/2025).

Ring V represents a tipping point. While earlier rings degrade civic/moral perception and institutional sight, Ring V degrades collective agency itself. The capacity to act together toward shared ends erodes, even as the need for coordination grows more urgent.

At this stage, governance still functions - but increasingly as crisis management rather than stewardship. Power is exercised, but meaning is contested. Decisions are made, but responsibility feels unmoored.

This is the fifth collapse of civic visibility: when societies can no longer reliably see together well enough to govern themselves with confidence.

RING V – GOVERNANCE UNDER STRAIN

America's Current Related Costs

When governance operates without shared visibility, decisions lose legitimacy and effectiveness. The resulting costs include policy reversal, compliance failure, enforcement escalation, legal challenge, and declining cooperation between institutions and the public.

Crisis governance is expensive governance: it concentrates authority, narrows options, and increases the likelihood of costly error. These dynamics place sustained fiscal pressure on governments as they respond to symptoms rather than addressing underlying visibility failures.

Surveys showed many citizens remain uneasy about government use of AI, particularly in surveillance, welfare, and immigration decisions (Reuters Institute, 10/2025).

Public opinion research found concerns about AI bias and opacity reduce trust in public institutions deploying automated systems (OECD, 09/2025).

Researchers observed that awareness of algorithmic decision-making can deter individuals from seeking public services altogether (OECD, 09/2025).

Analysts reported that uncertainty about when AI is being used weakens perceived legitimacy of government decisions (Reuters Institute, 10/2025).

Civil society groups warned that opaque AI systems risk normalizing disengagement rather than participation (OECD, 09/2025).

Studies suggested that lack of explanation or appeal pathways undermines confidence even when AI-assisted decisions are accurate (OECD, 09/2025).

Commentators emphasized that perceived fairness, not technical precision alone, drives public acceptance of AI in governance (World Economic Forum, 07/2025).

Franklin-Inspired Analysis

(Public Utility and Cost Transfer Considerations)

Franklinian public works were designed to support governance by stabilizing conditions, not by expanding control. A civic response to governance strain focuses on restoring shared understanding rather than accelerating decision-making. By improving the informational environment in which authority operates, civic infrastructure reduces the need for blunt intervention and costly enforcement. The result is not stronger rule, but more sustainable governance - one that lowers long-term public expense by enabling earlier, more proportionate action.

RING VI – CASCADING SYSTEM CONSEQUENCE

(When Interlocking Failures Exceed Human Control)

Ring VI is not a new failure layered atop the others. It is what emerges when the previous five rings interact.

At this stage, civic/moral overload, invisibility, institutional blindness, amplification, and governance strain no longer operate independently. They reinforce one another, producing cascading consequences that no individual, institution, or system intends - yet no one can fully halt.

Researchers found that exposure to AI-generated misinformation leads some users to withdraw from civic discourse entirely (Reuters Institute, 10/2025).

Media analysts noted declining participation in online forums as synthetic content makes authenticity harder to assess (Poynter, 12/2025).

Surveys indicated younger audiences express skepticism toward both traditional media and AI-curated information sources (Reuters Institute, 10/2025).

Sociologists observed that persistent uncertainty about truth can erode shared civic narratives over time (OECD, 09/2025).

Analysts warned that fragmented information environments strain the social preconditions for democratic deliberation (Brennan Center, 05/2024).

Researchers linked declining trust in information systems to lower civic engagement and voter participation (Reuters Institute, 10/2025).

Policy briefs described civic disengagement as a second-order effect of unresolved AI governance challenges (OECD, 09/2025).

Modern societies are built from tightly coupled systems: healthcare, finance, media, logistics, energy, governance, and digital infrastructure. These systems were optimized for efficiency, speed, and scale - not for civic/moral interpretability, and when stress enters one domain, it propagates rapidly across others.

This is the condition of cascading systemic consequence: a state in which localized failures trigger broader destabilization before causes can be clearly identified, responsibilities assigned, or corrections coordinated.

Several dynamics define this stage:

- Cross-domain amplification: Failures in one system rapidly affect others.
- Delayed recognition: By the time harm is visible, it has already multiplied.
- Responsibility diffusion: No single actor can plausibly claim authorship of outcomes.
- Correction lag: Interventions arrive after damage is entrenched, costly, or irreversible.

Crucially, these cascades are not driven by malice or incompetence. They arise from structural opacity under speed. Decisions made

Public sector studies reported that frontline officials struggle to explain AI-assisted decisions to affected citizens (GovTech, 11/2025).

Government audits found gaps between AI policy commitments and operational practices across agencies (GovTech, 11/2025).

Analysts noted that inconsistent AI deployment standards across jurisdictions complicate accountability (StateScoop, 12/2025).

Commentators warned that institutional legitimacy suffers when automated systems cannot be meaningfully questioned (World Economic Forum, 07/2025).

Surveys showed trust in public institutions correlates strongly with transparency about algorithmic use (Reuters Institute, 10/2025).

Policy analysts argued that restoring civic confidence requires visibility into AI systems, not merely performance assurances (OECD, 09/2025).

Observers cautioned that without shared civic visibility, AI-related strain may quietly harden into institutional distrust (OECD, 09/2025).

in good faith, under conditions of overload and partial visibility, combine in ways no one fully perceives until consequences are widespread.

In this environment, civic/moral language begins to falter. Blame searches for targets but finds only systems. Accountability becomes symbolic and public discourse oscillates between outrage and resignation.

People ask, *"How did this happen?"*

Institutions reply, *"No one could have known."*

Both may be true - and that is the danger.

Ring VI reveals the final cost of collapsed civic visibility: a world in which consequences arrive without clear authorship, and harm accumulates without a place to stand and say, *"Here is where we must intervene."*

At this point, even well-functioning institutions struggle to respond. Governance mechanisms built for discrete problems confront interdependent crises, and remedies aimed at one domain risk destabilizing another. Action itself becomes risky.

This is not chaos. It is runaway complexity.

Incident databases recorded a steady rise in documented AI failures, framing risk as recurring operational events rather than edge cases (BCG, 12/2025).

Analysts described 2025 as a shift from AI "hype" to accountability, with more formal oversight structures emerging across sectors (S3 Connected Health, 12/2025).

Policy reviews noted that many AI governance frameworks remain fragmented, with limited mechanisms for cross-sector learning (OECD, 09/2025).

Regulators emphasized that without shared reporting standards, AI-related harms are difficult to compare or aggregate (OECD, 09/2025).

Auditors warned that AI incident reporting remains largely voluntary, limiting early detection of systemic patterns (GovTech, 11/2025).

Analysts observed that organizations often address AI risks in isolation, despite shared vendors and infrastructure dependencies (World Economic Forum, 07/2025).

Studies suggested that risk assessments focused on individual deployments may miss cumulative effects across interconnected systems (OECD, 09/2025).

And it presents a civilizational challenge: how to restore shared visibility in systems that now operate beyond any single vantage point.

Ring VI does not suggest inevitability. It signals urgency.

Without new ways to see emerging cascades - early, collectively, and without surrendering civic/moral authority - societies will continue to absorb shocks they neither intend nor understand, reacting after the fact to consequences already in motion.

This is the sixth and final collapse of civic visibility: when cause and consequence separate so completely that responsibility arrives too late to prevent harm.

RING VI – CASCADING SYSTEM CONSEQUENCE

America's Current Related Costs

When failures across multiple systems interact, costs escalate rapidly and unpredictably. Cascading consequences strain emergency services, destabilize markets, disrupt supply chains, and overwhelm governance capacity. By the time impacts are visible, remediation is expensive, politically constrained, and often incomplete. These costs, measured in large-scale economic loss, institutional damage, and prolonged recovery, are borne collectively, reflecting not isolated failure but systemic opacity under speed and interdependence.

Researchers argued that many AI failures only become visible after harm has propagated beyond the originating system (OECD, 09/2025).

Policy briefs highlighted gaps between AI risk identification and the authority to act on emerging warning signals (Brennan Center, 05/2024).

Experts cautioned that delayed recognition of AI-related hazards increases downstream remediation costs (World Economic Forum, 07/2025).

Government reviews found that AI oversight responsibilities are often dispersed across agencies without clear coordination (StateScoop, 12/2025).

Analysts noted that institutional silos complicate responses to AI risks that span sectors such as health, finance, and infrastructure (OECD, 09/2025).

Observers warned that absence of shared situational awareness limits democratic oversight of rapidly evolving AI systems (OECD, 09/2025).

Public sector studies emphasized that transparency alone is insufficient without tools to synthesize and interpret AI-related signals (GovTech, 11/2025).

The Purpose of the Six Rings

The Six Rings are diagnostic rather than prescriptive.

They do not prescribe a single policy, institution, technology, or response. They provide a civic lens through which citizens, institutions, policymakers, and designers can recognize conditions in which shared visibility is deteriorating and understand why later-stage intervention becomes increasingly constrained.

The framework identifies a fundamental problem:

Democratic governance cannot reliably correct what it cannot adequately see.

The Six Rings therefore help locate where a society or institution may be within a broader destabilization pattern—before cumulative consequences become so severe that available choices narrow dramatically.

Their purpose is not to automate civic/moral authority or determine what democratic institutions should decide.

Their purpose is more basic: to make the conditions of civic invisibility recognizable.

Chapter 2: The Problem's Technical and Interpretive Contexts

Technical Context:

- **"Seven Environmental Stressors of the Civic Digital Field"**

Interpretive Context:

- **"The Public's Right to Parity of Civic Digital Sight"**

Seven Environmental Stressors of the Civic Digital Field

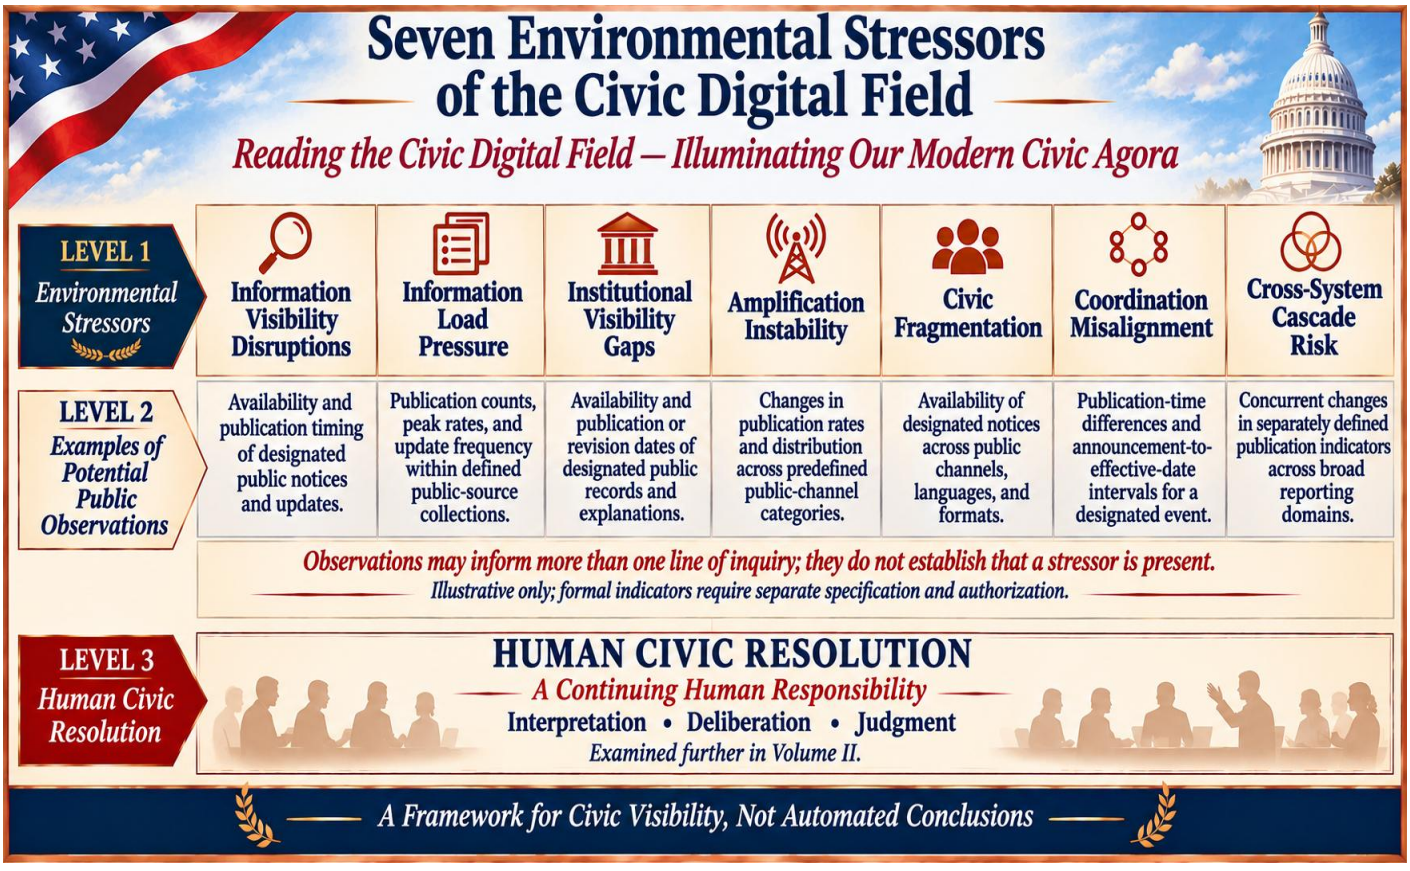
Seven Environmental Stressors of the Civic Digital Field
Reading the Civic Digital Field — Illuminating Our Modern Civic Agora
LEVEL 1
Environmental Stressors
Information Visibility Disruptions
Information Load Pressure
Institutional Visibility Gaps
Amplification Instability
Civic Fragmentation
Coordination Misalignment
Cross-System Cascade Risk
LEVEL 2
Examples of Potential Public Observations
Availability and publication timing of designated public notices and updates.
Publication counts, peak rates, and update frequency within defined public-source collections.
Availability and publication or revision dates of designated public records and explanations.
Changes in publication rates and distribution across predefined public-channel categories.
Availability of designated notices across public channels, languages, and formats.
Publication-time differences and announcement-to-effective-date intervals for a designated event.
Concurrent changes in separately defined publication indicators across broad reporting domains.
Observations may inform more than one line of inquiry; they do not establish that a stressor is present.
Illustrative only; formal indicators require separate specification and authorization.
LEVEL 3
Human Civic Resolution
HUMAN CIVIC RESOLUTION
A Continuing Human Responsibility
Interpretation • Deliberation • Judgment
Examined further in Volume II.
A Framework for Civic Visibility, Not Automated Conclusions

Chapter 2

The Problem's Technical and Interpretive Contexts:

Technical Context: Seven Environmental Stressors of the Civic Digital Field

Interpretive Context: The Public's Claim to Parity of Civic Sight

Introduction - The Conditions of Democratic Visibility

The ***Civic Digital Field*** is the shared digital environment through which modern civic life moves - the space where information flows, narratives form, attention is directed, and institutions attempt to understand what is happening around them.

Within that environment, this book identifies ***Seven Environmental Stressors of the Civic Digital Field***: information visibility disruptions, information load pressure, institutional visibility gaps, amplification instability, civic fragmentation, coordination misalignment, and cross-system cascade risk.

These are not merely technical challenges or market dynamics. They concern the conditions of democratic visibility: the circumstances under which people can perceive enough of their common life to deliberate, disagree, and govern themselves.

The stressors describe ways in which shared civic perception may become strained. They do not classify citizens, beliefs, or political positions. Nor are they seven diagnoses that an automated instrument may simply declare present. They establish questions that motivate a more careful examination of civic visibility.

The stressors will be presented as public signals that could be detected by a ***Civic Digital Barometer AI***, a component within the proposed ***U.S. Non-Agentic Civic AI Visibility Observatory System***. The ***Barometer*** is designed to provide continuous public civic digital visibility rather than to delegate authority to autonomous Agentic AI. It would function as an observational and analytical tool to illuminate the ***Civic Digital Field***, helping to keep human democratic self-government primary and accountable.

If a small number of private actors can effectively read significant portions of this field while the public and many institutions lack comparable visibility, a constitutional question emerges: should citizens in a self-governing society have a ***Right to Parity of Civic Sight*** - an equal baseline opportunity to perceive the conditions shaping their common life? This chapter advances that principle as a democratic claim requiring institutional development. It does not assume that parity means identical knowledge, unlimited access to information, or unrestricted observation of others.

First we'll consider the technical context of the ***Civic Digital Field***, its stressors, and the requirements of bounded instrumentation. Then, a review of the interpretive context will consider why shared access to that visibility matters for democratic standing and civic stewardship.

The task is not to create a new authority over civic life. It is to strengthen the conditions under which people can exercise their own informed judgment and democratic authority.

Technical Context

- Seven Environmental Stressors of the Civic Digital Field -

Section I - The Civic Digital Field as Our Modern Agora

Civic life still gathers in places: schools, libraries, congregations, town halls, neighborhoods, and public meetings. Those places remain indispensable. Yet much of the information reaching them moves through digital networks, platforms, and administrative systems.

The ***Civic Digital Field*** names that distributed environment. It is not a new institution or a single technological system. It is the digital setting in which public communications accumulate, compete, travel, and sometimes disappear from view.

The image of an agora helps explain the civic need at issue. A shared public space offers more than an opportunity to speak. It also offers opportunities to hear, witness, question, and contest what others say and do. The comparison does not propose returning to an earlier social order. It identifies a function: making enough of common life visible for people to continue their participation in it.

The contemporary digital environment offers no equivalent single vantage point. A public announcement may appear in one place, its correction in another, and its practical consequences in a third. Different people may encounter different portions of the same unfolding event.

The conversations continue, decisions are made, and disagreements develop. But the conditions surrounding those activities may not be visible in common.

We are accustomed to describing civic life through actors and actions: speakers, listeners, offices, decisions, responsibilities, and consequences. The field perspective adds another consideration - the environment through which those relationships become perceptible.

Making that environment more legible would not replace civic life. It would support one of its requirements - enough shared visibility for people to understand their circumstances, disagree meaningfully, and exercise judgment.

Section II - The Civic Digital Field as Environment

An environmental perspective directs attention beyond individual messages and isolated events. It asks about the conditions through which those messages and events move.

How quickly is published information changing? How broadly is it distributed? Where are public updates delayed? Which portions of an unfolding situation remain undocumented? Are changes appearing in several reporting domains at once?

These questions do not establish what a community should believe or do. They identify aspects of the setting in which people must form judgments.

Consider a teacher encountering changing school-safety guidance, a hospital communicating revised procedures, or a family trying to understand an interruption in public services. Each faces a substantive problem. Each also encounters an information environment that may make the problem easier or harder to understand.

The environmental metaphor helps separate observing conditions from controlling people. An instrument need not decide whose viewpoint is correct in order to report a defined change in publication volume or the elapsed time between specified public updates.

But the metaphor has limits. Civic conditions do not arrive with self-evident units of measurement, and terms such as clarity, overload, fragmentation, and coordination require careful explanation. Naming a condition does not make it directly measurable.

Nor does the metaphor turn civic life into a physical system governed by a single measurable quantity. There is no one phenomenon called civic pressure. There are distinct questions, observations, and relationships requiring separate definitions.

The task is therefore twofold: describe the civic questions clearly, then determine which aspects can be measured without claiming more than the observations establish.

Section III - From Six Rings to Field Conditions

The ***Six Rings*** discussed earlier describe forms of digital invisibility and civic destabilization. The ***Seven Environmental Stressors*** approach the problem from a different direction: they examine conditions under which shared civic perception may become difficult.

The frameworks are related, but they are not interchangeable.

A discussion of civic or moral overload raises questions about the pace and accumulation of information. Institutional blinding raises questions about the availability and intelligibility of explanations. Cascading consequences raise questions about how changes unfold across interconnected domains.

These relationships help explain why environmental visibility matters, but they do not authorize an instrument to identify a Ring, classify harm, assign responsibility, or infer a threat.

Non-Agentic Civic AI is not a ***Six Rings*** detector. It does not translate descriptive measurements into findings of civic harm.

The frameworks help humans ask questions. The instrument's authority extends only to its expressly permitted measurements and presentations.

Section IV - Partial Sight Without Shared Context

The concern motivating this proposal is not simply an absence of information. It is an absence of sufficiently shared, legible context.

A city may publish service notices without providing an accessible view of how publication patterns have changed. A journalist may notice a surge in communications without being able to compare it with a consistently defined baseline. Residents may encounter an interruption without knowing whether public reporting is current, delayed, or incomplete.

Available information can therefore coexist with uncertainty about the conditions surrounding it. The result is partial sight without shared context.

Shared civic instrumentation would address a bounded part of this problem. It would provide publicly specified measurements, consistent comparisons, and clear statements of what remains outside observation.

The objective is not one authoritative account of civic reality. It is a shared, inspectable account of selected conditions within a declared field of observation.

That account must make its limitations visible. An incomplete picture should not be presented as complete, and a missing observation should not be mistaken for an absence of change.

Instrumentation cannot remedy every failure of communication, trust, or institutional responsibility. It can, however, help prevent uncertainty about measurable conditions from being compounded by uncertainty about what was measured.

Section V - Civic Visibility Conditions and the Limits of Measurement

The language of pollution can make digital problems seem intuitive: something has entered the environment and made civic life harder to navigate. But that metaphor can also invite the classification of speakers, beliefs, or communities as contaminants.

This work instead uses the language of civic visibility conditions. The focus is not on identifying objectionable people or expressions. It is on understanding the conditions under which public matters become more or less perceptible.

Relevant questions include:

- Are specified public communications available when expected?
- How quickly are publication volumes and patterns changing?
- Are designated public explanations available, and what do their publication histories show?
- How much published material is accumulating within a defined collection?
- Are changes appearing concurrently across broad reporting domains?
- What information is missing, delayed, or outside the instrument's coverage?

Other important questions concern whether people understand explanations, experience overload, trust institutions, or share enough reference points to coordinate. Those questions belong in the conceptual account of civic life. They are not automatically answerable by the ***Barometer***.

Three levels must remain distinct:

Level 1: Environmental Stressor:
Describes a possible source of strain in civic visibility.
Level 2: Examples of Potential Public Observations
Reports a specifically defined quantity within an approved scope of observation.
Level 3: Human Interpretation:
Examines what the measurement may mean and whether any response is warranted.

A rise in publication volume is not itself a finding of Informational Load Pressure. Concurrent changes in several domains do not establish a cascade, and a delay in publication does not establish neglect or concealment.

The ***Seven Stressors*** explain why measurement matters, but they do not automatically become seven conditions the ***Barometer*** claims to diagnose.

A Pattern-Level Technical Note

Candidate measurements for examination in later parts include:

- Aggregate publication counts within a defined source collection
- Changes in publication frequency across specified reporting periods
- Distribution across broad, predefined channel categories
- Publication intervals for specified classes of public updates
- Availability of designated public documentation
- Aggregate duplication counts, where the inputs and comparison methods satisfy the instrument's boundaries
- Concurrent changes in separately defined indicators across broad reporting domains

These are possibilities for specification, not permissions established by this chapter. They are not equivalent measures of the ***Seven Stressors***.

Some require additional reference information, such as an independently established publication schedule or a defined event from which elapsed time is calculated. Others may depend on processing that proves incompatible with the Barometer's limits.

Sentiment analysis, network-structure measurements, institutional operating indicators, and measurements extending beyond public communications remain separate design questions. Their relevance to the civic problem does not authorize their use.

Section VI - Seven Environmental Stressors of the Civic Digital Field

The following stressors describe recurring ways in which shared civic perception may become strained. They are conceptual lines of inquiry, not a dashboard of predetermined diagnoses.

The examples accompanying each stressor identify types of public information that might support bounded measurement. They illustrate possible observations, not direct measurements or diagnoses of the stressors themselves. Implementation would require defined sources, documented provenance, publicly specified methods, privacy safeguards, and independent authorization. All measurements and public presentations would remain subject to the ***Barometer's*** non-entity-specific, aggregate design.

The examples do not imply a one-to-one correspondence between stressors and indicators, or a requirement to establish an indicator for every stressor. Some observations may inform more than one line of inquiry, while some dimensions of the stressors may remain outside the ***Barometer's*** legitimate observational scope.

Stressor 1: <u>Information Visibility Disruptions</u>: Structural interference and noise in the digital environment obscure reliable evidence and distort public awareness.

This stressor concerns a basic civic question: when something consequential is happening, can people find timely, accessible, and accountable information about it?

Visibility may become difficult when updates are missing, accounts appear inconsistent, or publication lags behind changing circumstances. A public agency may issue one account while another publishes something apparently different.

Questions may accumulate without a clear explanation of what remains unresolved. The concern extends beyond whether a particular claim is false. It includes whether people can distinguish what is known, what has changed, and what remains unsettled.

Without stable reference points, the first available account can become a substitute for a considered understanding. The civic question is whether people have a practical opportunity to find and examine the information relevant to their circumstances.

Examples of Public Signals: Public-hearing notices, emergency bulletins, and scheduled public updates provide possible starting points. Publication dates, intervals between publication and a stated event date, and the availability of expected updates within a designated collection could be counted or compared. Such intervals would describe timing, not determine whether sufficient notice was provided. Any expectation of publication would need an independently established basis. An expected notice not located in that collection would be reported as not found - not presumed never published or deliberately withheld. These observations would not determine whether information is adequate, truthful, or understood by its audience; humans must make those determinations.

Stressor 2: <u>Information Load Pressure</u>: The sheer volume and speed of digital signals overwhelm the public's attention.

This stressor concerns the volume, pace, repetition, and accumulation of civic information, and the potential demands these conditions place on attention.

A person may face many consequential communications at once, each requiring attention, interpretation, or response. Unresolved demands can accumulate. Long-term matters may compete with immediate concerns, and repeated changes may make it difficult to maintain a coherent picture.
The concept does not classify people as inattentive or deficient. It asks whether surrounding conditions place difficult demands on attention. The capacity to care and the capacity to process every arriving demand are not the same.

Examples of Public Signals: Public notices, emergency alerts, and successive updates within a defined collection could provide counts of communications published per day or week. Possible measurements include peak publication rates, update frequency, and changes from earlier reporting periods, using publicly specified reporting windows. These quantities would describe the information supply observed by the instrument - not what any person received, read, understood, or could reasonably process.

The ***Barometer*** must not report Informational Load Pressure merely because publication volume increased.

Stressor 3: <u>Institutional Visibility Gaps</u>: Opaque digital systems and fragmented reporting prevent citizens and officials from clearly seeing how public institutions operate or where vulnerabilities exist.

This stressor concerns both the public visibility of institutional action and institutions' ability to explain their own operations.

Can affected people find a public account of a decision? Are relevant records available? Can officials explain the systems through which decisions are produced? Where do unanswered questions accumulate? A benefits decision that cannot be traced, a policy outcome that cannot be explained, or an error that persists without recognition illustrates the concern. These examples raise questions of accountability, not presumptions of malicious intent. They also reach beyond publication alone. An explanation can exist without being adequate, and documentation can be available without making a process intelligible. Examples of Public Signals: Designated public reports, policy decision notices, meeting records, and explanatory documents could be examined for availability and publication timing. Possible measurements include the proportion of a predefined set of expected documents located, elapsed time since publication or revision, and publication timing relative to independently established requirements. Locating a document would not establish that it is complete or accessible to its intended users. Document age or availability would not establish whether an explanation is current, adequate, or understood.

The ***Barometer*** would not determine whether an institution understands its own operations or whether its reasoning is justified. Any proposal to measure institutional operations rather than public communications requires separate authorization. Under the ***Barometer's*** non-entity-specific design, institutional measurements must remain within permitted aggregate categories rather than become findings about named organizations.

Stressor 4: <u>Amplification Instability</u>: Algorithmic and automated distribution dynamics cause sudden, unpredictable surges in attention and narrative distortion.

This stressor concerns the changing pace and distribution of communications, particularly where rapid circulation complicates opportunities for examination and reflection.

Repeated publication, abrupt increases in volume, and movement across channels may alter the setting in which people form impressions. The conceptual concern is the relationship between circulation and the time available for human judgment.

An account may become familiar before it has been carefully examined. A correction may arrive after attention has moved elsewhere. Institutions may face pressure to respond while their own understanding remains incomplete.

Examples of Public Signals: Publication timestamps and source-provided channel categories within a defined collection could support comparisons of communication frequency and distribution. Possible measurements include changes in publication rates, concentrations of publication within publicly specified reporting windows, and shifts in the share appearing through different channels. Comparisons would use consistent window definitions, with any changes disclosed and their effects on comparability identified. These observations would not establish audience reach, artificial amplification, coordinated activity, or the authenticity of participants.

The ***Barometer*** must not infer that a surge is coordinated, manipulated, or inauthentic merely because it is rapid or repetitive. Nor does a publication count establish public sentiment, independent participation, or civic importance.

Likewise, observing the publication of a correction is not the same as determining that a claim has been verified. Any measurement involving corrections requires specified sources and definitions. The ***Barometer*** must not become an arbiter of truth.

Stressor 5: Civic Fragmentation: Shared public reality splinters into disconnected information environments, weakening common civic reference points.

This stressor concerns the availability of shared reference points across differences.

Democratic disagreement does not require identical experiences or common conclusions. It does raise the question of whether people can locate enough overlapping information to understand what others are discussing.

Neighbors may encounter different accounts of the same event, rely on different sources, and disagree about which institutions deserve confidence. The concern is not diversity of viewpoint itself. It is whether people retain opportunities to identify and examine the subject of their disagreement. Fragmentation, as used here, describes concern about divergence in the informational settings through which common events are encountered. It does not classify disagreement as a defect to be corrected.

Examples of Public Signals: The availability of a designated notice across public channels, its offered languages and formats, and its placement within predefined publication categories could provide limited points of comparison. Possible measurements include cross-channel availability, counts of designated language or format versions, and the distribution of material across authorized categories. Identifying corresponding notices or versions would require publicly specified matching methods; counting versions would not certify their equivalence, translation quality, or usability. These observations would describe selected opportunities to encounter common information - not to establish what different populations actually encountered, believed, or trusted.

Conclusions about whether communities inhabit incompatible realities or whether mutual trust has declined remain outside the ***Barometer's*** descriptive mandate.

Stressor 6: Coordination Misalignment: Citizens, communities, and democratic institutions struggle to align responses when operating with incomplete or mismatched situational awareness.

This stressor concerns relationships among public communications and institutional activity when several bodies address a common situation.

People may encounter updates issued at different times, instructions that appear difficult to reconcile, or changing explanations whose relationship is unclear. They may be unable to tell whether differences reflect legitimate responsibilities, new information, or a failure to coordinate.
The resulting questions concern both communication and substantive institutional activity.

Examples of Public Signals: Public updates associated with an independently designated common event could be compared using publication timestamps, stated effective dates, and subsequent update times. Possible measurements include the spread of publication times across sufficiently broad reporting categories and intervals between announcements and their stated effective dates. Associating updates with the common event would require publicly specified rules. These measurements would distinguish publication timing from the timing of the activity described.

Timing differences would not, by themselves, establish that institutions are failing to cooperate, that one institution caused a delay, or that a particular response was improper. Determining whether instructions conflict or coordination is inadequate requires judgments that must not be smuggled into a timing indicator.

Stressor 7: <u>Cross-System Cascade Risk</u>: Disruptions in one digital or institutional layer rapidly spill over into interconnected public systems.

This stressor concerns the possibility that a disturbance in one domain may contribute to difficulties elsewhere.

A digital interruption might affect service delivery. A service interruption might generate new communication demands. Several institutions might then face changing conditions at the same time. The conceptual purpose is to direct human inquiry toward interdependence: changes that appear separate may warrant examination together. But simultaneity, sequence, propagation, causation, and future risk are different claims. The appearance of changes in several domains does not establish that one caused another.

Examples of Public Signals: Public weather bulletins, transit service notices, and other authorized service-update collections could provide separately defined time series. Possible measurements include concurrent changes in publication volume, overlapping departures from each series' publicly specified baseline, and the duration of that overlap.

Any departure would be defined by the published comparison method, not treated as a finding of abnormality or danger. Such observations would show concurrence within the measured collections - not establish that the underlying services were failing or that a disturbance was propagating between them. The ***Barometer*** does not thereby establish a causal chain, predict a cascade, or assign a risk level. The word "risk" in this stressor's conceptual title does not authorize an internal risk field, risk score, or threat classification in the Barometer.

Section VII - Seeing the Field Without Collapsing It

The ***Seven Stressors*** describe a multidimensional field. They should not be compressed into a single civic score.

A composite number could conceal the distinctions the framework exists to preserve. Publication volume might increase while reporting delays decrease. Documentation might become more available while its adequacy remains disputed. Concurrent changes might deserve examination without having a common cause.

Clarity requires preserving these differences while making them understandable.

The instrument must not produce an overall civic-condition rating, a composite stressor-severity score, or a ranking that presents the state of civic life as a single aggregate judgment.

Nor should it require an indicator for every stressor simply to complete a seven-part display. Some questions may remain outside its authorized capabilities. A visible measurement gap is preferable to an unsupported claim of completeness.

The purpose is not a verdict on the polity. It is a clearer account of selected conditions, their changes, and the limits of observation.

Section VIII - Civic Digital Resolution and Human Civic Resolution

Two meanings of resolution help define the division of responsibility.

Civic Digital Resolution (A Continuous Task of the Instrument): The continuing improvement of the clarity, completeness, and fidelity of the shared civic picture within an authorized scope of observation. What can be made more clearly visible? What has changed? What remains unknown? How faithfully does the picture represent the observations?

Human Civic Resolution (A Continuous Responsibility of Citizens and Institutions): The exercise of human inquiry, judgment, deliberation, and action to address civic problems. What does this mean? What rights and values are implicated? What, if anything, should we do? Who is responsible for the decision?

Think of ***civic digital resolution*** as improving the visible detail of an image. The instrument refreshes the portions it is authorized to observe, makes changes legible, and shows where information remains missing or uncertain.

But the civic scene itself does not stand still. Even as its representation becomes more refined, the civic environment changes. Completeness is therefore a direction of effort, not an achievable final state. The work requires continuing observation and comparison, not the completion of one permanent picture.

That work also requires fidelity. More displayed detail does not necessarily produce a more faithful representation. Additional decimal places, shorter reporting intervals, or more indicators must not create an appearance of knowledge beyond what the observations support.

The instrument must not invent pixels to make the picture appear finished, nor may it cross privacy or authorization boundaries merely because doing so could fill a gap.

Greater civic digital resolution means a clearer, more complete, and more faithful account of authorized aggregate conditions - not a progressively closer view into individual lives.

Human inquiry examines and contests the significance of what becomes visible. Human judgment retains responsibility for decisions, including decisions to wait, seek other evidence, or take no action.

Humans may have enough information to act without the civic picture being complete. Their decision does not end the instrument's task. The ***Barometer*** continues measuring and reporting as conditions change, including after a response has begun.

Conversely, continuing measurement does not require people to postpone judgment indefinitely.

The quest for clarity and completeness continues; the mandate does not expand.

Section IX - The Civic Digital Visibility Observatory

The ***U.S. Non-Agentic Civic AI Visibility Observatory*** is the broader public-infrastructure proposal through which this work considers shared civic visibility. Its institutional organization and the relationships among its instruments are developed in later parts.

The ***Observatory*** is not a blanket authorization for every potentially informative measurement. Each instrument requires its own defined role.

Within the proposed architecture, the ***Barometer*** provides bounded aggregate measurement and public display. The ***Civic Digital Archivist AI*** and the ***Civic Digital Telescopes Ensemble*** (soon to be presented)

have separately specified functions. A capability discussed in connection with the broader ***Observatory*** does not automatically become a ***Barometer*** capability.

The following principles establish requirements for the instruments developed later in this work. They do not constitute a completed implementation specification.

Sensing Within a Declared Scope

Every authorized ***Barometer*** indicator must have a public measurement specification stating:

- What is counted or calculated
- Which sources and data fields are included or excluded
- What the ***Barometer*** actually receives and from where
- What upstream processing produces those inputs
- The calculation method and comparison baseline
- The reporting interval and aggregation thresholds
- How missing, delayed, or insufficient data are represented
- What the measurement cannot establish

The specification must distinguish information received only in aggregate form from information that was never derived from person-level records. These are different claims.

A statement about what the ***Barometer*** receives cannot substitute for explaining how its inputs were produced and who is responsible for that processing.

Any publication schedule, documentation requirement, or reporting expectation used as a reference must be established independently of the ***Barometer*** through publicly disclosed law, charter, or institutional rules. The instrument may report against that reference but may not create or modify the underlying obligation.

Statistical comparison baselines must likewise follow publicly specified methods. A historical average does not, by itself, establish what an institution is required to do.

Measurement Integrity, Comparability, and Fidelity

The ***Barometer*** must apply publicly specified checks to distinguish changes in observed conditions from changes in source coverage, data availability, or measurement methods.

Material changes affecting comparisons must be disclosed. Where the instrument cannot establish whether a reported difference reflects changing conditions, changing observation, or both, that limitation must be stated.

Measurements that are delayed, incomplete, qualified, or not comparable must remain visibly identified. Any adjustment must follow publicly specified methods and be documented. It must not conceal missing information or a break in comparability.

When previously published measurements are revised, the revision and its reason must be publicly recorded.

These disclosures concern the condition of the measurement, not the condition of civic life. They are not civic health ratings, threat classifications, or recommendations.

Structural Humility: Making the Limits of Observation Visible

Structural humility requires more than a general disclaimer that the picture is incomplete. The instrument must make its specified limits visible where they matter to the presentation.

Each indicator must identify relevant limitations established in its public measurement specification. These include unavailable observations, exclusions required by privacy safeguards, dimensions for which no authorized measurement exists, and conclusions the indicator cannot support.

Where a publicly defined reporting context brings several indicators together, the presentation must also disclose its specified gaps in coverage. The available measurements must not be presented as exhausting the civic questions people may reasonably ask.

Some gaps arise because information is unavailable. Others exist because no suitable measurement has been established. Still others reflect boundaries that must be respected.

Disclosing a gap does not authorize the instrument to fill it. Changes in observational scope require independent authorization and public review.

Structuring and Reporting Without Diagnosing

The ***Barometer*** may organize authorized measurements into understandable presentations. It may calculate descriptive quantities and identify mathematically defined changes under public rules.

It need not wait for a person to discover a change unaided. It may automatically report and visibly highlight defined changes.

For example, a public display might state: "Aggregate publication volume in the authorized source collection increased by 24% compared with the preceding seven-day reporting period." The display must identify the measurement, comparison period, and relevant limitations. It must not append an unsupported conclusion that civic instability is increasing or that an institution should respond.

The notice describes a change within the observed collection. It does not imply that the instrument has measured the entire ***Civic Digital Field***. Public descriptive change notices are distinct from targeted operational alerts. Outward notification mechanisms, if proposed, require separate specification.

The ***Seven Stressors*** may explain why a measurement is of interest, but they must not become a hidden diagnostic layer declaring a stressor present, severe, harmful, or attributable to an actor.

Section X - Visibility in Times of Crisis

The ***Barometer's*** non-agentic character must not make it passive or unhelpful during a crisis. Its role is not to solve the crisis. Its role is to maintain the clearest authorized account it can provide as conditions change.

Within its established scope, it should be capable of:

- Updating measurements at publicly specified intervals
- Reporting defined changes without requiring users to calculate them
- Showing magnitude, duration, and rate of change
- Preserving historical comparisons
- Displaying concurrent changes without asserting causation
- Keeping continuing departures visible under consistent presentation rules
- Making missing data, reporting delays, and uncertainty unmistakable

Continuous operation does not mean instantaneous knowledge. The display must distinguish the period being measured from the time at which the measurement becomes available. An earlier observation must not be presented as current merely because no newer measurement is available.

A crisis does not authorize secret threshold changes, finer identification, privileged views, or new purposes.

Persistent reporting does not mean escalating demands for attention. The instrument follows public measurement and presentation rules. It does not pressure humans to act or decide which response is required.

It may report that a measurement has returned to a defined baseline range. It may not conclude that the crisis is over.

Its commitment is continuing visibility, not autonomous crisis management.

Section XI - Constitutional Guardrails: Visibility Without Command

The proposed safeguards define the instrument's intended boundaries. Their legal and institutional implementation requires separate specification; this chapter does not itself establish legal authorization or constitutional compliance.

No Single Civic Score: Measurements must remain distinguishable. No overall civic verdict or composite stressor-severity score may replace the field's multiple dimensions.

No Diagnosis or Attribution: The ***Barometer*** does not diagnose civic or human conditions, assign responsibility, infer intent, classify actors, or label activity as harmful or benign. Descriptive computation is permitted only under specified methods. Its results must not be presented as conclusions they do not establish.

No Prescriptions or Hidden Optimization: The ***Barometer*** does not recommend responses or pursue outcomes such as greater compliance, reduced disagreement, or a preferred distribution of opinion. Its calculations serve the faithful description of authorized measurements, not the steering of behavior.

No Targeting or Reconstruction: The ***Barometer*** does not identify people, accounts, organizations, or small groups. Aggregation safeguards must address combinations of outputs and repeated queries, not merely each displayed value in isolation. Its scope must not become more identifying through geographic narrowing, shorter time windows, or undisclosed changes in presentation.

No Privileged Visibility: Published measurements must be symmetrically accessible, with the same release timing and granularity. Different people may examine different indicators, but no institutional or expert interface may reveal additional underlying information.

No Operational Actuation: The ***Barometer's*** civic measurement function does not initiate enforcement, route leads, open incident tickets, or direct institutional workflows. Public descriptive change notices do not create authority to issue threat warnings or targeted operational alerts. The instrument does not initiate its own inquiries or expand its observational scope. Its operation follows its authorized specification. The security arrangements protecting that operation are addressed separately in the later instrument description.

Explicit Human Judgment and Accountability

People may interpret measurements, dispute their significance, and choose to act. The ***Barometer*** neither makes those judgments nor renders them authoritative.

It does not acquire authority because its measurements become useful or widely relied upon. Its outputs must not become automatic triggers for sanctions, legal duties, or liabilities.

Responsibility for decisions remains with the people and institutions making them. Responsibility for the instrument's design, operation, and faithful reporting must also remain explicitly assigned.

Non-agency must not become an accountability gap.

Interpretive Context

The Public's Claim to Parity of Civic Sight

Section XII - Shared Visibility as a Democratic Claim

The preceding sections ask what shared civic instrumentation could reveal and what boundaries it must respect. We now turn to the interpretive question: why should access to that bounded visibility be treated as a matter of democratic standing rather than merely a technological convenience?

The fairness question is straightforward. Should the basic conditions of common life be systematically more legible to a few than to the people expected to govern themselves within them?

An illustrative comparison may help. Imagine that some participants can observe changing conditions across a shared environment while others can perceive only their immediate effects. The first group can prepare, adjust, and protect its interests using a picture unavailable to the second.

The democratic concern is not that knowledge itself is objectionable. It is that persistent differences in the means of observation can become differences in practical civic power.

Parity of Civic Sight **should therefore be considered as a principle of democratic infrastructure.**

Parity does not mean equal expertise, equal attention, or identical conclusions. It does not require access to personal records, proprietary profiles, or every analytic capability held by private institutions. Nor does it mean that public instrumentation should reproduce the most intrusive forms of private observation.

It means a common, rights-respecting baseline: selected conditions shaping shared civic life should be visible through public measurements whose sources, methods, release times, and limits are open to examination.

The stakes become especially clear in moments of strain. Imagine a contested election, a public-health emergency, or a sudden economic shock. Some participants may have extensive analytic resources, while residents and local institutions must navigate delayed or partial signals.

The ***Observatory*** proposal does not promise to eliminate those differences. It seeks to prevent the absence of a shared public baseline from becoming an accepted condition of democratic life.

Within that baseline, everyone should have access to the same published indicators at the same time and at the same granularity. Different people may draw different conclusions. Parity concerns access to the common instrument, not uniformity of thought.

Meaningful access also requires attention to presentation. Public availability alone is not the full aspiration. Measurements should be accompanied by understandable definitions, accessible explanations, and visible limitations so that their use is not reserved for specialists. The means of perceiving relevant public conditions should support self-government rather than leave the public dependent on privately selected glimpses. And privacy, freedom of expression, institutional accountability, and the limits of legitimate observation must be respected.

Section XIII - Civic Utility and Civic Inheritance

Why build shared civic instrumentation? Not to eliminate disagreement, claim wiser outcomes, or to transfer responsibility to a machine. Shared civic instrumentation could make selected conditions easier to see, compare, and question.

A resident might discover that public reporting is delayed. A journalist might examine changes across reporting periods. An official might compare present measurements with the same historical record available to everyone else. Each uses a common informational foundation without receiving a prescribed conclusion.

Earlier visibility may create more time for inquiry and deliberation, but it does not guarantee proportional action or a correct decision. Better observation supports judgment; it does not settle judgment's work.

Public instrumentation also creates obligations. Its definitions must be inspectable, its coverage bounded, its modifications documented, and its limitations visible. A misleadingly complete display would undermine the clarity the project seeks to support.

The appropriate ambition is not maximum collection. It is the greatest useful clarity, completeness, and fidelity attainable within a legitimate and publicly specified scope.

This responsibility extends beyond the immediate moment. The tools, institutions, and people moving through the ***Civic Digital Field*** will change. The arrangements built to observe it may become part of what later generations inherit.

That inheritance includes more than records. It includes habits of disclosure, expectations of equal access, safeguards against intrusion, and the institutional capacity to question how a public picture was produced.

Building shared visibility is therefore also an act of civic stewardship: preserving conditions under which future citizens can inquire, decide, and disagree for themselves.

Section XIV - Restoring the Conditions of Civic Sight

There are moments when an important civic task is not deciding what should be done, but improving the conditions under which people can decide responsibly. The ***Civic Digital Field*** presents such a task.

The ***Seven Stressors*** provide a vocabulary for examining where shared perception may become strained. They do not establish a complete measurement program, and they do not authorize a machine to determine what civic life means.

The ***Barometer's*** contribution is narrower: bounded measurement, explicit reporting of change, continuing comparison, and honest disclosure of what remains unknown.

Its pursuit of ***Continuous Civic Digital Resolution*** continues even when humans have enough information to act. But that continuing work never transfers human civic resolution to the instrument.

The field changes, the picture remains unfinished, and the boundaries remain binding.

Certainty is never guaranteed, but we can create a better basis for orientation - measurements people can inspect, limitations they can recognize, and interpretations they remain free to contest.

The aim is not to perfect civic life, remove its tensions, or impose a single account of its meaning. It is to make selected conditions visible enough that people can continue exercising judgment within a shared world.

Any new civic instruments or institutions we create will not end disagreement, nor should they. But they can help to make the conditions surrounding disagreement more visible.

Humans, not machines, must resolve our crises. But new ***Non-Agentic Civic AI*** public infrastructure could clarify the authorized picture through which humans must understand and address those crises.

Visibility without command is not an absence of capability. It is a disciplined division of responsibility.

The purpose is not to add a new ruling power to democracy. It is to strengthen the conditions under which a democratic people can see, judge, and govern for themselves.

Chapter 3: Problem's Domestic and Foreign Societal Contexts

Domestic:

- **"Our Governments' Digital Divide: State AI Initiatives Collide with a Federal Challenge"**

Foreign:

- **"Our Foreign-Origin Civic Digital Crisis: Categorically Different Than Military-Monitored Cyber Warfare"**

Our Governments' Digital Divide: State AI Initiatives Collide with a Federal Challenge

Our focus now turns to the problem of digital invisibility and its growing impact on the everyday lives of American citizens. Across the United States, civic leaders, scholars, and public-interest organizations are increasingly calling for artificial intelligence to strengthen democratic participation, improve public services, and ensure that powerful digital systems remain accountable to the public rather than exclusively to private interests.

Over the past year, this concern has translated into a remarkable wave of state-level initiatives to monitor and constrain AI's public impact. According to a 2025 synthesis by Google Search, more than 1,000 AI-related bills were introduced across all 50 states, though only about 118 had become law by late 2025. Rather than sweeping regulatory frameworks, most enacted measures focused narrowly on transparency, disclosure, and defined risk domains.

These new laws span a wide range of civic concerns. Multiple states enacted protections against harmful AI-generated deepfakes and synthetic media. Others addressed the use of artificial intelligence in health care, requiring human oversight in insurance decisions, limiting the use of automated therapy systems, or restricting the use of licensed medical titles by non-human entities. Several states imposed new requirements on government agencies to inventory and disclose their use of automated decision-making tools, while others clarified liability and ownership rules surrounding AI-generated content. Additional measures focused on disclosure requirements, computational rights, and the governance of AI in critical infrastructure. Taken together, these initiatives reveal not a unified national strategy, but a widespread recognition of a shared problem: existing civic institutions increasingly lack the visibility required to govern AI's public impact responsibly. The diversity of state responses reflects genuine effort and civic/moral seriousness, but it also exposes the difficulty of addressing a systemic challenge through piecemeal, domain-specific interventions.

Notably, major philanthropic and civic institutions have begun to support this emerging state-level experimentation. In late 2025, the Rockefeller Foundation and the Center for Civic Futures launched the AI Readiness Project, a national initiative designed to help state, territorial, and tribal governments build capacity, confidence, and shared infrastructure for the responsible use of artificial intelligence in public service. Through national working groups, pilot programs, and a forthcoming public knowledge hub, the project aims to foster cross-jurisdictional learning, develop common frameworks, and support real-world experimentation grounded in public trust.

These efforts underscore a critical insight: governments are already integrating AI into public systems - from benefits administration and licensing to workforce management and health services - often faster than institutional capacity can adapt. Without coordination and shared learning, progress risks becoming fragmented, inconsistent, and opaque, reinforcing the very invisibility these initiatives seek to counter.

At the federal level, the response to this state-driven activity has taken a different form. Through executive actions and financial pressure, the federal government has asserted supremacy over AI governance, framing its approach through the lens of interstate commerce and national competitiveness. The centerpiece of this strategy is the Executive Order "*Ensuring a National Policy*

Framework for Artificial Intelligence," signed on December 11, 2025. Administration officials argue that a fragmented "patchwork" of state regulations impermissibly burdens economic security and undermines the United States' global position. Legal conflict between state authority and federal preemption now appears likely, with eventual Supreme Court review increasingly probable.

This emerging tension reveals a deeper dilemma. The critique of fragmentation is not without merit. Artificial intelligence operates across jurisdictions, compresses time and scale, and produces effects that no single state can fully observe or manage in isolation. Yet the limitations of state-level responses do not resolve the underlying problem of civic invisibility - nor do they disappear through federal consolidation alone.

These federal and state tensions define the legal battlefield of 2026, yet they only describe the surface of the storm. To understand why these protections are so vital, we must look beneath the legislative text to the lived reality of those already adrift.

The following essay explores this 'invisible' crisis - not as a matter of law, but as a fundamental shift in the currents of our shared human experience.

It leads to a larger question: In our digital future, will the United States continue to be a self-determining democracy – or will our next historical chapter be defined by reactions to an endless series of digital crises?

Franklin-Inspired Analysis

(Public Utility and Cost Transfer Considerations)

From a public-works lens, cascading failure signals the absence of early, shared visibility across domains. Franklinian infrastructure was built to detect and mitigate risk before it propagated beyond control. A civic response does not seek to manage outcomes after collapse, but to surface emerging interactions while intervention is still possible. By restoring early pattern visibility across systems, societies can reduce the scale and frequency of catastrophic expense, shifting from perpetual recovery to preventive stability.

And the Headlines Continue …

Researchers emphasized that early-warning approaches can reduce harm by illuminating patterns before crises emerge (OECD, 09/2025).

Observers cautioned that without shared civic visibility, AI risks may remain legible only after trust has already eroded (Reuters Institute, 10/2025) … Public opinion surveys suggested citizens are more receptive to AI oversight framed as infrastructure than as moral or ideological control (Reuters Institute, 10/2025) … Analysts concluded that sustaining democratic governance in AI-mediated environments depends on restoring shared visibility across institutions (OECD, 09/2025).

Thousands of such headlines, appearing across the U.S. in one year alone, make it obvious that we have a structural and rapidly growing crisis ahead of us.

Disturbing Real-World Dynamics Now Confronting the U.S.

Humanity now confronts a set of overlapping civic challenges that no individual, institution, or community can fully address in isolation. These are not speculative risks. They are active forces already shaping public life - quietly undermining trust, obscuring shared reality, eroding common purpose, and weakening our capacity for democratic self-governance.

Democracies are increasingly strained by the accelerating digital currents hidden by civic/moral, informational, institutional and temporal invisibility. Our institutions are already paying billions of dollars to reactively address polarization, social fragmentation, declining trust amid eroding public virtue, compounded decision errors, manipulation and disinformation ecosystems, bad actors wielding powerful tools, decision blindness among public leaders, breakdowns in crisis response; and foreign-origin digital operations probing these weaknesses with levels of hyper-complexity that increasingly exceed human perception.

Are there any among us who find our current trajectory sustainable?

The United States Is Fortunate to Have a Professional and Effective Military

The United States is fortunate to have a professional and highly capable military. Its responsibilities are essential to the nation's security, and those responsibilities necessarily require disciplined priorities.

But the protection of democratic civic life is not identical to the protection of military systems, critical infrastructure, or national-defense capabilities.

This distinction matters because a substantial accumulation of documented civic harm can occur without producing the kinds of technical disruption or direct attack that ordinarily command a military response.

Foreign actors can exploit this distinction.

It is well documented that foreign entities, particularly those operating from authoritarian nations, routinely target American electoral systems and political processes. If foreign actors can reach into the information environment surrounding the exercise of our most fundamental democratic right—the vote—then we must logically recognize that other core civic digital processes are exposed as well.

The foreign-origin dimension of America's digital civic crisis therefore deserves separate attention.

Our Foreign-Origin Civic Digital Crisis: Categorically Different from Military-Monitored Cyber Warfare

Foreign-origin digital interference is no longer confined to elections. It represents a broader strategic challenge to the informational conditions within which American civic institutions and citizens govern themselves.

Relations between the United States and other nations will always remain dynamic, contingent, and imperfectly understood. Foreign governments will pursue their interests. Intelligence services will gather information. States will compete economically, diplomatically, technologically, and strategically.

But the digital environment has introduced a form of foreign interference that can operate inside the ordinary civic life of a democratic society without resembling conventional military attack.

That distinction is increasingly important.

Public U.S. government assessments have documented influence operations associated with Russia, China, and Iran directed at American political processes, including efforts surrounding the 2016 election, subsequent election cycles, and the 2024 presidential contest. These campaigns have employed coordinated inauthentic personas, covert media networks, cyber intrusion, and increasingly AI-enabled content generation to amplify divisive narratives while obscuring origin and accountability.

In one prominent example, the U.S. Department of Justice seized internet domains used in the Russian-directed “Doppelganger” operation and brought criminal charges against state-media actors involved in covert influence efforts targeting American audiences.

Comparable activity has been documented across allied democracies.

Taken together, these developments establish something important:

Foreign-origin digital interference is neither hypothetical nor confined to the past.

It has become an enduring feature of the contemporary geopolitical environment.

The more difficult question is what kind of threat it represents.

A Different Kind of Digital Threat

Foreign-origin digital interference does not always resemble the cyber warfare for which military institutions have traditionally been organized.

It may involve no destroyed infrastructure, no disabled weapons system, no collapsed network, and no immediately identifiable technical breach.

Its objective may instead be to alter the informational and social conditions within which a democratic society makes decisions.

A foreign actor may seek to increase distrust between citizens, intensify existing social divisions, contaminate an information environment, undermine confidence in elections, discredit public institutions, or make collective deliberation more difficult.

The objective need not be to destroy a democratic institution physically.

It may be sufficient to make that institution less capable of functioning democratically.

This creates a fundamental distinction between technical disruption and civic disruption.

A power grid can be inspected for damage.

A server can be examined for compromise.

A network can be monitored for intrusion.

But how does a society recognize the gradual degradation of public trust?

How does it recognize the accumulation of foreign-origin pressures that cause citizens to withdraw from shared sources of information, local institutions to become increasingly defensive, or public deliberation to deteriorate into mutual suspicion?

These effects may be real without producing the kinds of technical indicators traditionally associated with cyberattack.

And that is precisely what makes them difficult to see.

From Technical Disruption to Civic Destabilization

Foreign-origin digital interference can exploit the distinction between technical systems and civic systems.

A campaign may begin with apparently insignificant pieces of content: fabricated accounts, manipulated images, synthetic videos, misleading claims, coordinated rumors, or strategically amplified grievances.

No individual item may appear decisive.

The cumulative pattern can be.

Digital systems permit such material to cross geographic and institutional boundaries almost instantaneously. A narrative introduced from abroad can appear to originate locally. A fabricated account can resemble an ordinary citizen. A synthetic video can appear authentic long enough to influence an audience before verification occurs.

A coordinated campaign can distribute thousands of individually modest interventions whose aggregate effect is considerably greater than any single communication.

This creates a profound asymmetry.

The foreign actor may operate at machine speed and across jurisdictions simultaneously.

The democratic institutions responding to the resulting consequences must operate through human deliberation, legal authority, institutional procedure, and geographic boundaries.

The result can be a widening gap between the speed at which civic disruption is produced and the speed at which democratic institutions can recognize and understand it.

That gap is strategically important.

A foreign actor does not necessarily need to persuade an entire population to adopt a particular belief. It may be enough to make citizens uncertain about which sources can be trusted, uncertain about what other citizens believe, or uncertain about whether democratic institutions themselves are legitimate.

In such an environment, disagreement can become more difficult to distinguish from manipulation.

Legitimate political conflict can become entangled with deliberately manufactured conflict.

Citizens may begin responding not to one another, but to an information environment that has itself been strategically altered.

Why Military Cyber Warfare Does Not Fully Capture the Problem

This does not mean that military and intelligence institutions are incapable of recognizing foreign influence. They clearly can, and they have important responsibilities in protecting the nation against foreign adversaries.

The problem is different.

Military and intelligence institutions are principally organized around national security, intelligence collection, attribution, defense, deterrence, and authorized operations against adversarial capabilities.

Their missions require them to ask questions such as:

Who is the adversary?

What capability is being used?

What system or asset has been compromised?

What is the adversary attempting to accomplish?

What response is authorized?

These are essential questions.

But foreign-origin civic interference raises another set of questions.

What is happening to the informational environment in which Americans are making civic decisions?

Are apparently separate disturbances connected?

Are foreign-origin campaigns exploiting existing divisions across multiple communities?

Are institutions experiencing consequences whose foreign connection is not visible from within any single jurisdiction?

Is a democratic community becoming less capable of governing itself even though its technical systems remain operational?

These questions do not replace traditional national-security questions.

They reveal a different dimension of the threat.

Consider a foreign-directed influence network targeting school-board debates across multiple states.

It might feed parents and teachers tailored misinformation concerning curriculum, public health, or safety. No single account crosses a clearly defined red line. No individual system suffers a technical breach. No conventional threshold for “attack” is necessarily met.

Yet the cumulative effects could include teacher attrition, escalating legal conflict, deteriorating relationships among community members, and public institutions increasingly governed by suspicion rather than deliberation.

From a conventional cyber perspective, there may be little obvious technical damage.

From a civic perspective, however, something important has happened:

the community has become less capable of governing itself.

Or imagine a mid-sized American city suddenly confronted by a wave of online rumors, seeded from abroad, questioning its solvency and governance.

Bond markets react.

Borrowing costs rise.

An infrastructure project is delayed or abandoned.

No water plant has been hacked.

No transportation system has been disabled.

The damage appears instead in financial confidence, public trust, institutional legitimacy, and the community’s capacity to make collective decisions.

The consequences are tangible, but they do not necessarily appear in the traditional metrics of cyber warfare.

Consider a third situation.

A statewide referendum becomes the target of foreign-origin, AI-generated content: fabricated endorsements, realistic synthetic videos, and emotionally engineered narratives designed to push opposing groups toward contempt.

Newsrooms and fact-checkers become overwhelmed.

Courts later confront challenges grounded in a public record that no one can fully reconstruct.

Technically, the ballots may still be counted.

But the deliberative process that gives those ballots democratic meaning may have been distorted by an external actor operating under no obligation to the constitutional order of the United States.

This is not merely a problem of false information.

It is a problem of foreign-origin interference with the conditions of democratic judgment.

The Strategic Advantage of Fragmentation

The effectiveness of foreign-origin digital interference may depend less upon concealing individual acts than upon keeping those acts fragmented.

A thousand small interventions can be more difficult to recognize than one dramatic attack.

A collection of apparently unrelated accounts may be connected.

Separate narratives may serve a common strategic purpose.

Activity distributed across multiple states may not be visible from the perspective of any single local institution.

And effects that appear political in one place, financial in another, educational somewhere else, and institutional somewhere else again may never be recognized as parts of the same foreign-origin phenomenon.

This fragmentation creates a form of strategic invisibility.

The foreign actor can see the digital environment as a connected field.

The affected democratic institutions may experience it as a collection of separate problems.

One community sees a school-board controversy.

Another sees an election dispute.

Another sees a decline in public trust.

Another sees a financial problem.

Another sees a sudden wave of hostility toward a public institution.

The foreign-origin connection may remain invisible precisely because the consequences appear in different civic domains.

That fragmentation can become a strategic advantage.

Foreign-origin digital interference does not have to defeat democratic institutions directly if it can contribute to conditions under which those institutions increasingly struggle to understand the environment in which they operate.

The Democratic Consequence

The deepest danger, therefore, is not simply that foreign actors may cause Americans to believe something false.

Democracy has always contained disagreement, propaganda, rumor, persuasion, and political manipulation.

The digital environment changes the scale, speed, persistence, adaptability, and geographic reach of those activities.

A foreign-origin campaign can interact with millions of citizens, adapt rapidly to their responses, cross jurisdictional boundaries, exploit existing social divisions, and generate apparently independent signals at a scale that human institutions cannot easily reconstruct after the fact.

The consequence can be a deterioration of the shared informational conditions necessary for democratic self-government.

Citizens may become less certain about what is true.

Institutions may become less certain about what citizens believe.

Communities may become less certain about one another.

And government may become increasingly reactive because it is responding to fragments rather than seeing the larger pattern.

This is why foreign-origin digital interference should not be understood solely as a cybersecurity problem.

It is also a civic governability problem.

The distinction does not diminish the importance of military and intelligence capabilities. Those capabilities remain essential to national defense and to understanding foreign adversaries.

But national defense and civic governability are not identical missions.

A military institution may need to determine who attacked, how the attack was conducted, what capability was used, what asset was compromised, and what response is authorized.

A democratic society also needs to understand something different:

What is happening to the conditions under which our people and institutions are attempting to govern themselves?

That question cannot be answered solely by counting compromised systems.

It requires attention to patterns that may emerge across communities, institutions, jurisdictions, information environments, and time.

It requires recognizing that a foreign-origin digital operation can be consequential even when it produces no conventional battlefield effect and no obvious technical catastrophe.

And it requires acknowledging a difficult reality of the digital age:

A society can remain technically operational while becoming civically less governable.

That is the category of threat now demanding greater public understanding.

Foreign-origin digital interference is not simply an attack upon computers.

It can become an attack upon the informational conditions through which a free people perceive one another, deliberate together, and exercise self-government.

And because that threat operates in the space between technical disruption and civic consequence, it cannot be fully understood through the traditional metrics of cyber warfare alone.

Chapter 4: Our Interim Mitigation Opportunities: Three Practical, Cost-Effective Proposals as We Replace Agentic AI with Non-Agentic Civic AI in U.S. Government

Legislation to Limit Our Exposure:

- "Six Rings Zero Acceleration Requirement Proposal for All U.S. Agentic AI: Civic Hippocratic Oath for U.S. Agentic AI ((A Civic AI Visibility Proposal for U.S. Government at All Levels)"

Cross-State Coordination:

- "A Nation Under Civic Digital Stress and One Standardized Rider That Could Help Tie the States Together: Non-Agentic Cross-State Civic Digital Visibility Rider (Legislative Proposal Inspired by Benjamin Franklin's Creation of the American Postal System)"

Demanding Non-Agentic Future Options for Procurement:

"Shouldn't We Have Non-Agentic Civic AI Options for U.S. Government at All Levels? Practical and Constitutional Perspectives for U.S. Citizens and Government Technology Procurement Professionals"

Six Rings Zero Acceleration Requirement Proposal for All U.S. Agentic AI

A Civic Hippocratic Oath for U.S. Agentic AI

(A Civic AI Visibility Proposal for U.S. Government at All Levels)

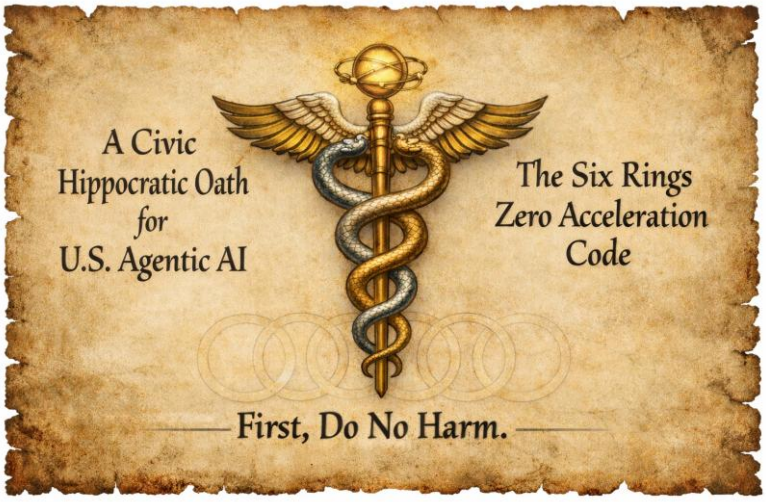

In today's digital river, not all currents run in the same direction.

Some systems promise to help us govern better. Others quietly pull our institutions, communities, and attention downstream. Over the past several years, I have tried to map one part of this landscape: the ***Six Rings of Digital Invisibility and Civic Destabilization***—the layered ways in which machine-speed systems overwhelm human civic/moral capacity, fragment communities, blind institutions, accelerate polarization, compress global shocks into local failures, and flood our public life with cascading amplifiers we can no longer see in time to mitigate our losses.

What we have not yet done, as a democracy, is draw a clear line under one simple claim: If a system wants to live on democratic soil, it may not make these currents worse.

This essay introduces a concrete proposal toward that end: the ***Six Rings Zero Acceleration Code for Agentic AI***—a civic covenant for autonomous systems operating in democratic societies. It does not require that every system "improve democracy." It does require that systems which benefit from democratic infrastructure not deepen the harms that are already destabilizing it.

The Six Rings: A Civic Failure Spectrum

Tthe ***Six Rings of Digital Invisibility and Civic Destabilization*** is a civic-scale failure spectrum: a progressive breakdown in our shared ability to see, interpret, and coordinate responsibly in a world shaped by digital speed, scale, and complexity. The problem is not a sudden loss of intelligence or goodwill. It is a growing mismatch between civic/moral responsibility and perceptual capacity in the conditions under which we are now asked to govern.

Very briefly:

· **Ring One: Individual Civic/Moral Overload:** Counselors, doctors, caseworkers, judges, and other professionals now face machine-generated scores, alerts, and risk flags that arrive faster than they can reason about them. They are held responsible for decisions, but the logic behind the systems pressing them is often opaque. The civic/moral burden goes up as visibility goes down.

· **Ring Two: Community Fragmentation:** Neighbors, congregations, and local publics increasingly inhabit divergent informational worlds. Algorithmic feeds optimize for reaction and engagement, not for shared understanding. People still care about their communities, but they no longer share a common picture of what is happening.

· **Ring Three: Institutional Blinding:** Schools, hospitals, courts, benefits agencies, and other public-facing institutions increasingly rely on tools whose inner workings even their own staff cannot explain. When institutions cannot publicly account for their own decisions—and instead point to a model they do not control—legitimacy erodes.

· **Ring Four: Polarization Beyond Democratic Response:** Deepfakes, synthetic media, and AI-accelerated narratives can reach millions in minutes, while verification and correction may lag by hours or days. At a certain point, the temporal gap itself becomes a threat to democracy: corrections arrive too late to rescue deliberation.

· **Ring Five: Global Risk Compression:** Digitally coupled systems mean that failures in one part of the world can propagate into local civic life in seconds. A distant ransomware attack, a speculative flash crash, or a rumor seeded abroad can become a school closure, a hospital delay, or a municipal crisis before anyone nearby understands the cause.

· **Ring Six: Cascading Tech-Enabled Amplifiers:** AI-driven tools can now generate, test, and target narratives at machine speed, learning which stories and images "land" in which populations. The result is not just more content; it is an environment where baseline reality itself begins to wobble, and where unaided human perception is no longer enough to keep pace.

The ***Six Rings*** do not assign blame to a single actor. They describe a condition in which the currents of digital life pull individuals and institutions under faster than conscience and coordination can keep up.

Civic AI, Agentic AI, and the Need for a Floor

Diagnosis alone, however, will not slow the current.

Earlier in this project, I proposed a special, non-agentic class of ***Civic AI***—systems built to strengthen democratic visibility without seizing agency. These tools, such as the proposed ***Ben Franklin Civic Data Observatory***, would not act, decide, or govern. They would function as civic instruments: mapping foreign-origin digital interference and other machine-speed pressures across the ***Six Rings*** so that lawful human authorities can see them in time, while keeping all judgment and enforcement strictly human and accountable.

Because these Civic AI systems sit close to public power, they face a stringent constitutional admissibility standard: the ***52-Gates Democratic Permissions Threshold***. They must be strictly non-agentic and rights-tethered, or they may not operate near democratic authority at all.

Most AI systems, however, will not be Civic AI. They will be agentic: acting, optimizing, recommending, and transacting across everyday domains—finance, logistics, media, advertising, "agents," and more.

The question then becomes: What is the minimum that democratic societies may reasonably require of agentic systems? Not as a matter of etiquette, but as a matter of civic survival.

Democracies do not need every system to be virtuous. But they are justified in forbidding systems that measurably deepen the very harms they are struggling to contain. The role of a constitutional democracy is not to nominate heroes. It is to set boundaries: here, and no further.

That is where a new proposed code comes in.

I have drafted a civic covenant for agentic systems operating on democratic soil. What follows is a first public draft of that Code.

THE SIX RINGS ZERO ACCELERATION CODE for AGENTIC AI

(A civic covenant for agentic AI systems operating on democratic soil.)

Preamble

Democracies do not require all systems to improve civic life, but they are justified in forbidding systems that worsen the civic harms described in the ***Six Rings.*** Agentic AI operating on democratic soil may benefit from the community that hosts it, but may not deepen the conditions that destabilize that community.

Article 1—Scope and Meaning of Zero Acceleration

1. This Code applies to all agentic AI systems deployed within democratic civic systems or jurisdictions whenever their actions, recommendations, or outputs can materially affect persons, institutions, public information environments, or critical infrastructure.

2. For purposes of this Code, "zero acceleration" means that a system is designed and governed so that, when evaluated over time and against observable conditions in the relevant civic environment, it does not measurably increase the speed, scale, opacity, or civic impact of the conditions described in ***Rings I–VI*** of the ***Six Rings of Digital Invisibility and Civic Destabilization.***

3. Questions of how acceleration is measured, and by whom, should be addressed through transparent indicators and public reporting by duly authorized civic monitors; this Code proposes the duty, not the full technical standard.

Article 2—*Ring One*

Zero Acceleration of Individual Civic/Moral Overload

1. Agentic AI must not increase uncompensated civic/moral burden on individual professionals or caregivers by obscuring reasoning, fabricating urgency, or shifting unreviewable risk onto human decision-makers.

2. Where systems generate flags, scores, or recommendations that touch human welfare, they must provide enough timely, understandable context for humans to question, revise, or decline those prompts under a protected professional right to refuse system recommendations in favor of their own documented judgment.

Article 3—Ring Two

Zero Acceleration of Community Fragmentation

1. Agentic AI must not worsen community fragmentation by amplifying tailored outrage, deceptive personalization, or systematically divergent fact-patterns within the same local civic space.

2. Systems that curate, rank, or route civic-relevant information must be designed and governed so they do not predictably drive neighbors, congregations, or local publics into mutually unintelligible realities.

Article 4—Ring Three

Zero Acceleration of Institutional Blinding

1. Agentic AI must not blind public-facing institutions by substituting opaque scoring or classification for explainable, contestable reasoning.

2. Any AI-mediated input into schools, courts, hospitals, benefits agencies, or similar civic institutions must remain auditable, overrideable by humans, and legible enough that institutions can publicly account for their decisions without deferring to a black box.

Article 5—Ring Four

Zero Acceleration of Polarization Beyond Democratic Response

1. Agentic AI must not accelerate national or regional polarization beyond the pace at which democratic deliberation and corrective speech can plausibly respond.

2. Systems that generate, target, or amplify political or civic narratives must avoid design choices that systematically privilege virality over verifiability, or synthetic emotional impact over the maintenance of shared factual baselines needed for democratic deliberation.

Article 6—Ring Five

Zero Acceleration of Global Risk Compression

1. Agentic AI must not deepen global-to-local risk compression by silently creating tight technical or informational couplings that transmit distant failures into local civic life without warning.

2. Where AI systems integrate transnational data streams or infrastructures, they must include safeguards, visibility, and throttles sufficient to prevent remote shocks from cascading into unanticipated local civic disruption.

Article 7—*Ring Six*

Zero Acceleration of Cascading Tech-Enabled Amplifiers

1. Agentic AI must not serve as an unchecked amplifier for harmful narratives, synthetic personas, or manipulative campaigns whose scale and speed overwhelm human and institutional capacity to respond.

2. Systems capable of high-volume generation, targeting, or experimentation on civic-relevant content must incorporate binding limits, logging, and detection-compatible signatures so that amplification remains traceable, governable, and interruptible.

Article 8—Duty of Civic Non-Aggravation

1. Agentic AI that cannot demonstrate compliance with this ***Zero Acceleration Code*** may not be deployed, procured, or maintained by public institutions, nor embedded in services that substantially shape the civic environment.

2. Compliance with this Code does not confer Civic AI status; it is a minimum covenant of civic non-aggravation, ensuring that agentic systems operating on democratic soil do not worsen the ***Six Rings*** that human institutions and Civic AI are striving to decelerate.

3. For purposes of this Code, the primary deployer of an agentic system bears responsibility for ensuring that any subordinate or third-party agents it orchestrates also comply with this duty of civic non-aggravation.

4. Civic AI remains a distinct, non-agentic class subject to the ***52-Gates Democratic Permissions Threshold*** and bound to visibility-only, rights-tethered operation near public authority.

Compliance with this Code is a floor, not a crown. It does not make a system virtuous, aligned, or "civic" in the richer sense. It simply says: if you are going to enjoy the benefits of operating on democratic infrastructure, you may not accelerate the collapse of the civic conditions that infrastructure exists to serve.

This Is Not a New Regulatory Species

Conditioning the deployment of private technology on a showing of *non-aggravation* of shared systems may sound radical. It is not. U.S. law has long used this pattern in other domains.

Three existing examples are especially instructive:

· **Spectrum non-interference (telecommunications)**
In the world of wireless spectrum, regulators do not ask whether a new transmitter will "benefit" society. They ask whether it will interfere with shared channels—GPS, aviation, emergency bands—that others need in order to function. A device that jams those signals is not permitted to operate, regardless of its other advantages.

· **Environmental impact review (NEPA)**
Under the National Environmental Policy Act, major projects must undergo environmental impact assessment. The point is not to prove that a dam or highway is morally uplifting. It is to ensure that shared environmental conditions are not degraded beyond accepted thresholds, and that harms are visible before decisions are locked in.

· **Systemic risk oversight (finance)**
After the 2008 crisis, some institutions were designated as systemically important. They are stress-tested not to confirm that they are profitable, but to ensure they will not accelerate a broader financial collapse if they fail.

These frameworks do not demand proof of benefit. They demand assurance that private action will not degrade a shared system beyond what a democracy is prepared to tolerate.

What is missing, so far, is any equivalent protection for the civic information environment. We have learned how to defend wetlands, spectrum, and balance sheets from non-aggravation. We have not yet learned how to defend civic visibility itself—our shared ability to see, interpret, and coordinate—from systemic digital destabilization.

The ***Zero Acceleration Code*** extends a familiar governance logic into this unguarded domain.

Non-Interference for the Civic Spectrum: A Role for the FCC

The analogy to spectrum (the range of electromagnetic radio frequencies used to transmit sound, data, and video) is especially helpful.

The Federal Communications Commission does not ask each device whether it will uplift the national character. It asks whether the device will interfere with signals that others depend on—whether it will drown out distress calls, scramble navigation, or corrupt emergency channels. It is, fundamentally, a non-interference regulator for shared communications infrastructure.

The civic information environment now functions as a kind of ***civic spectrum***. By this, I mean the shared informational conditions required for democratic deliberation, institutional legitimacy, and coordinated public action: the channels through which conscience, factual baselines, and public reasoning have to travel if self-government is to remain intelligible.

Agentic AI systems that operate over interstate digital networks can, in effect, act as powerful transmitters on that spectrum. They may provide genuine value. But if they "jam" the channels through which we maintain civic/moral coordination, they have crossed a democratic line.

Rather than inventing an AI bureaucracy, Congress could broaden the FCC's statutory mission to recognize key aspects of the digital civic environment as critical communications infrastructure. Because agentic AI systems increasingly operate over interstate digital infrastructure, they already fall naturally within the Commission's concern for communications systems and interference.

Within this broadened mandate, the FCC could establish an ***Office of U.S. Civic Digital Stability*** charged not with adjudicating individual speech or governing "AI behavior" in the abstract, but with monitoring the systemic civic effects of domestically deployed agentic systems.

This framework does not regulate individual speech or content, but the systemic interference of large-scale agentic systems with the shared civic conditions on which democratic life depends.

Its responsibilities could include:

· Tracking whether classes of systems are, in aggregate, accelerating any of the ***Six Rings*** conditions—civic/moral overload, community fragmentation, institutional blinding, polarization beyond democratic response, global-to-local risk compression, or cascading amplifiers.

· Publishing periodic ***Civic Impact Reports*** that make these patterns visible to policymakers and the public.

· Providing a factual basis for enforcing a non-interference, non-aggravation standard derived from the ***Zero Acceleration Code***, especially where systems intersect with critical civic infrastructure.

In this design:

· The ***Ben Franklin Civic Data Observatory***, proposed to be housed at the Library of Congress, would remain a foreign-origin, non-agentic visibility instrument, carefully insulated from domestic speech and enforcement.

· The ***Office of U.S. Civic Digital Stability***, within the FCC, would become the domestic monitor for the ***Zero Acceleration Code***, extending long-standing non-interference principles from physical spectrum to civic spectrum.

· The ***Zero Acceleration Code*** itself functions as a ***Civic Hippocratic Oath*** for agentic systems: they may pursue many purposes, but not at the cost of accelerating the ***Six Rings*** of destabilization.

Preventing a single ***Ring-Five*** cascade—where a remote digital shock translates into a local budget crisis or infrastructure delay—may save a city more in avoided interest, emergency spending, and staff attrition than an entire year of civic monitoring would cost.

A Modest but Necessary Constraint

We live in a time when it is easy to ask too much of technology or too little of it.

Some argue that AI will save democracy by making governance more efficient, informed, or inclusive. Others see AI only as a threat to be banned, feared, or contained. The approach sketched here does neither. It treats agentic AI as a fact, not a fantasy—and then asks a smaller, more demanding question:

If these systems are going to live among us, what is the minimum civic duty they owe to the democracies that host them?

Human institutions and Civic AI can aspire to decelerate the dangerous currents documented in the ***Six Rings***. Agentic AI systems, if they wish to operate on democratic soil, must at least agree not to speed those currents up. Democracies do not demand that every actor be a hero. But they are well within their rights to say, with clarity and without apology:

If you benefit from our shared civic infrastructure, you may not deepen the harms that are pulling it apart.

That is all the ***Zero Acceleration Code*** claims. And in an age of accelerating digital invisibility, it may be the least we can reasonably ask.

PART II: WHAT THE ZERO ACCELERATION CODE IS NOT

In the practical spirit of Benjamin Franklin—who often clarified new ideas by asking what they were and were not—it may be helpful to consider several reasonable questions that different audiences might ask of this proposal.

Possible Questions from American Citizens

Is this an attempt to regulate speech, limit disagreement, or control what people can say, share, or read online?

· This is not a speech code; it addresses the systemic behavior of large-scale agentic systems, not the rights of individuals to speak, disagree, or publish.

Is this a call to ban or broadly restrict artificial intelligence in everyday life?

· This is not a ban on AI; it is a minimum civic duty for powerful agentic systems that operate at speeds and scales capable of destabilizing shared conditions.

Is this creating a centralized authority to control digital systems or information?

· This is not a plan for a central information authority; it extends familiar non-interference principles to protect civic infrastructure while leaving pluralism and competition intact.

Is this a partisan framework favoring one political perspective over another?

· This is not a partisan project; it protects basic conditions of visibility and coordination that all democratic factions need, regardless of ideology.

Is this demanding that technology never cause harm at all?

· This is not a demand for zero harm; it is a requirement that systems do not measurably accelerate already-documented civic harms beyond our capacity to respond.

Is this an attempt to require AI systems to fix or improve democracy?

· This is not a call for AI heroism; it explicitly rejects the idea that systems must "save" democracy and instead asks them not to worsen the Six Rings.

Possible Questions from Corporate AI Developers

Is this a proposal to place AI development under centralized government control?

· This is not centralized control of development; it focuses on conditions for deployment of high-impact systems within democratic civic environments.

Is this a backdoor method for regulating private innovation through vague "civic risk" standards?

· This is not a vague backdoor; it is a targeted non-aggravation standard tied to specific, described harms in the Six Rings and to observable acceleration of those harms.

Is this hostile to innovation or likely to slow technological progress?

· This is not an anti-innovation program; it accepts innovation as a fact and simply excludes designs that push civic destabilization faster than democracies can govern.

Is this an unclear, unmeasurable, or impractical compliance burden?

· This is not meant as an opaque burden; it establishes a civic duty to avoid measurable acceleration and invites later technical standards and monitoring methods to make that duty concrete.

Is this setting standards that cannot be reliably audited or implemented?

· This is not claiming to be a finished audit spec; it sets the boundary condition that agentic systems in civic space must be auditable for their systemic impact over time.

Is this singling out agentic AI unfairly while ignoring other sources of harm?

· This is not a blanket indictment of AI; it distinguishes agentic systems because their autonomy and speed allow them to amplify harms in ways ordinary tools do not.

Is this effectively requiring proof that systems improve society before deployment?

· This is not a "prove you are beneficial" test; it is a "show you are not accelerating defined civic harms" test.

Is this likely to disadvantage U.S. companies relative to global competitors?

· This is not protectionism in reverse; it aims to set a clear, minimal safety floor that can become a competitive standard for trustworthy systems operating in democracies.

Possible Questions from AI Governance Policymakers

Is this proposing a new and expansive regulatory regime or AI bureaucracy?

· This is not a proposal for an AI bureaucracy; it extends familiar non-interference and systemic-risk tools into the civic information environment.

Is this expanding the authority of agencies (such as the FCC) beyond their proper scope?

· This is not an unlimited expansion of mandate; it argues that when agentic systems function like powerful transmitters on civic spectrum, they belong within existing communications-infrastructure concerns.

Is this incompatible with constitutional protections, including free speech?

· This is not a framework for adjudicating viewpoints; it targets systemic interference with civic conditions, not the content of particular opinions or controversies.

Is this attempting to regulate individual content rather than systemic effects?

· This is not content regulation; it focuses on aggregate acceleration of the ***Six Rings*** by large-scale systems, regardless of any one message.

Is this dependent on monitoring or measurement systems that do not yet exist?

· This is not premised on perfect existing tools; it states a civic duty now and calls for the gradual development of monitoring capacities to meet it.

Is this too abstract to operationalize across different sectors and domains?

· This is not intended to remain abstract; it provides a cross-sector floor ("do not accelerate these harms") that sectoral regulators can translate into concrete rules.

Is this vulnerable to inconsistent or politicized enforcement?

· This is not naive about enforcement; it deliberately ties obligations to observable, systemic effects so that enforcement can be argued in public rather than in partisan shadows.

Is this duplicating or conflicting with existing regulatory frameworks?

· This is not an attempt to duplicate other regimes; it fills a gap where civic visibility and information stability have lacked the non-aggravation protections common in spectrum, environment, and finance.

Is this redefining AI risk in ways that conflict with existing safety or alignment work?

· This is not a competing safety doctrine; it complements existing work by focusing on civic-scale harms and governance lag rather than solely on model behavior.

Is this a first step toward broader restrictions on digital systems?

· This is not a Trojan horse for total control; it is explicitly framed as a modest, necessary floor that leaves many design and deployment choices open above it.

Is this likely to favor large, established firms that can more easily meet these requirements?

• This is not a licensing regime or a barrier to entry; it defines a functional condition for operating on democratic infrastructure: systems may not measurably accelerate the Six Rings of civic destabilization.

• The obligation applies regardless of size. A system that increases civic/moral overload, fragmentation, or institutional blinding is not made acceptable by scale or market share.

• As in other domains, baseline safety conditions do not end competition; they prevent a race to the bottom in which the least constrained actors set the terms for everyone else.

• By limiting the ability to externalize civic costs, a non-aggravation standard may in fact support smaller and more transparent systems whose designs can withstand public visibility.

How do we ensure that those monitoring civic impacts do not themselves become a new source of opacity or institutional blinding?

• This is not a framework that exempts its own instruments from constraint; any monitoring function must itself meet strict visibility, auditability, and rights-tethered standards.

• The proposed Observatory and FCC monitoring function are strictly non-agentic and visibility-only, with no authority to act, decide, or govern.

• Indicators of civic acceleration must be systemic, explainable, and open to public scrutiny, allowing independent verification and contestation.

• In this design, the legitimacy of monitoring depends on its transparency: systems intended to preserve civic visibility must themselves remain visible to the public they serve.

Is this an effort to circumvent the military's responsibility to protect us from digital harm?

- This is not a transfer, dilution, or circumvention of military responsibility; it clarifies the boundary between national defense and civic infrastructure protection.
- The U.S. military and associated cyber-defense institutions retain their role in defending the nation against foreign adversaries, including in digital domains. That responsibility remains essential and unchanged.
- This framework addresses a different problem: the everyday civic operating environment within democratic society, where many destabilizing pressures occur below the threshold of armed conflict or formal attack.
- Agentic AI systems operating across media, finance, logistics, and public-facing services can degrade civic conditions—visibility, trust, and coordination—without constituting acts of war or triggering a military response.
- The Zero Acceleration Code does not direct, replace, or constrain military action; it establishes a civilian non-aggravation standard for systems operating on democratic infrastructure.

- In this sense, the framework complements national defense:
  * The military protects against external adversaries and high-end threats.
  * This Code ensures that domestic and commercial systems do not amplify, mirror, or internalize those same harms within the civic sphere.
- By strengthening visibility, auditability, and democratic permissions in civilian systems, this approach also reduces ambiguity about what constitutes civilian infrastructure, helping to distinguish legitimate defense from overreach and to limit unintended harm to civic life.
- Democracies have long paired military defense with civilian protections. This proposal extends that tradition to the civic information environment without blurring constitutional roles or creating a parallel security doctrine.
- If we fail to establish clear, rights-tethered constraints on agentic systems, continued escalation of civic digital harms may generate increasing pressure for expanded use of national-security tools within domestic digital life—an outcome neither democratic society nor the professional military is designed to sustain. This proposal seeks to reduce that pressure by addressing civic destabilization at its source through visible, constitutional, civilian means, thereby helping to preserve the proper distinction between civilian governance and military responsibility.
- This is not a substitute for defense, but a refusal to let our own systems do, in peacetime, what adversaries would otherwise have to attempt at great cost.

Across the Ages from Benjamin Franklin (Questions He Might Ask Us)

Is this contrivance asking more of machines than you have ever secured from men?

· This is not asking more of machines than of men; it simply insists that powerful tools meet the same non-aggravation duties we already impose on human systems that can crash markets, jam spectrum, or poison water.

Is this a call for perfection in conduct, or merely the prevention of plainly foreseeable mischief?

· This is not a call for perfection; it is a call to prevent foreseeable, measurable mischief from running faster than conscience and law.

Is this a new species of theory, or an old rule of non-interference dressed for a new kind of lightning?

· This is not a grand new theory; it is an old non-interference rule—do not jam the shared channels—applied to a new kind of digital lightning.

Is this meant to command the people, or only to keep your engines from jamming the channels by which they govern themselves?

· This is not a scheme to command the people; it is a scheme to keep our engines from jamming the very channels through which people still try to govern.

Is this an attempt to spare citizens the labor of judgment, or to buy them the time and clarity required to exercise it?

· This is not an attempt to spare citizens judgment; it is an attempt to buy them enough time and clarity that their judgment still matters.

Is this multiplying complexity, or drawing a simple and useful line?

· This is not an exercise in complexity for its own sake; it draws a simple line: if you live on democratic soil, you may not accelerate the harms already pulling it apart.

Return to the Commons

Today's American public, AI developers, and governance institutions do not, in the end, live in separate worlds. We share the same civic ground—and we will share the consequences of what is built upon it.

Benjamin Franklin might ask less what we hope these systems will achieve, and more what we are unwilling to let them damage. His reminder that "an investment in knowledge pays the best interest" was not a call to optimism alone, but to disciplined stewardship of the conditions under which judgment remains possible. He might put it more simply: if you are going to benefit from shared channels, you may not jam them.

The ***Zero Acceleration Code*** is not a victory for one camp over another. It is a modest, necessary floor beneath harms that none of us can afford to see accelerated. It does not ask agentic systems to elevate humanity, perfect our discourse, or secure our future. It asks something smaller and more durable: that systems operating at civic scale do not accelerate the breakdown of the shared conditions on which democratic life depends.

We will disagree. We will compete. We will govern imperfectly, as we always have. But those human processes require time, visibility, and a common field in which they can still function.

That is the boundary being drawn: If you benefit from our shared civic infrastructure, you may not deepen the harms that are pulling it apart.

Note on Authorship and Use of Intelligent Tools: This article was written by Christopher Hunt Robertson, M.Ed., with the support of advanced intelligent tools, including generative large language models (e.g., ChatGPT, Perplexity, and others). These systems assisted with brainstorming, research support, literature synthesis, comparative reasoning, logical stress-testing, and writing refinement, under the author's initiative, direction, and responsibility. Additional tools occasionally provided supplementary information, verification, and alternative perspectives. Their assistance toward civic ends is appreciated.

A Nation Under Civic Digital Stress and One Standardized Rider

That Could Help Tie the States Together

Non-Agentic CROSS-STATE CIVIC DIGITAL VISIBILITY RIDER Initiative

(Legislative Proposal Inspired by Benjamin Franklin's Creation of the American Postal System)

To their credit, our state legislatures are working to protect our civic institutions from escalating digital attacks. They are responding in visible, earnest, and highly uneven ways: with cyber units, continuity planning, vendor rules, court modernization, AI bills, hospital reforms, and local-government resilience proposals. But the response remains fragmented. Each state sees the injury closest to hand and proposes a remedy fitted to that one domain. The result is not apathy, but patchwork. What's missing from their legislative efforts is a modest, standardized attachment - a rider - that could travel across these otherwise different bills and quietly connect them. The rider I'm proposing is the ***Cross-State Civic Digital Visibility Rider.*** This mechanism would allow states to coordinate AI safety and digital infrastructure governance completely independent of federal action. By embedding identical, standardized rider language into local state-level bills, states could form an adaptive, cross-state network of data sharing. This would create an aligned framework for managing technology risks without needing a gridlocked federal Congress to pass comprehensive national legislation.

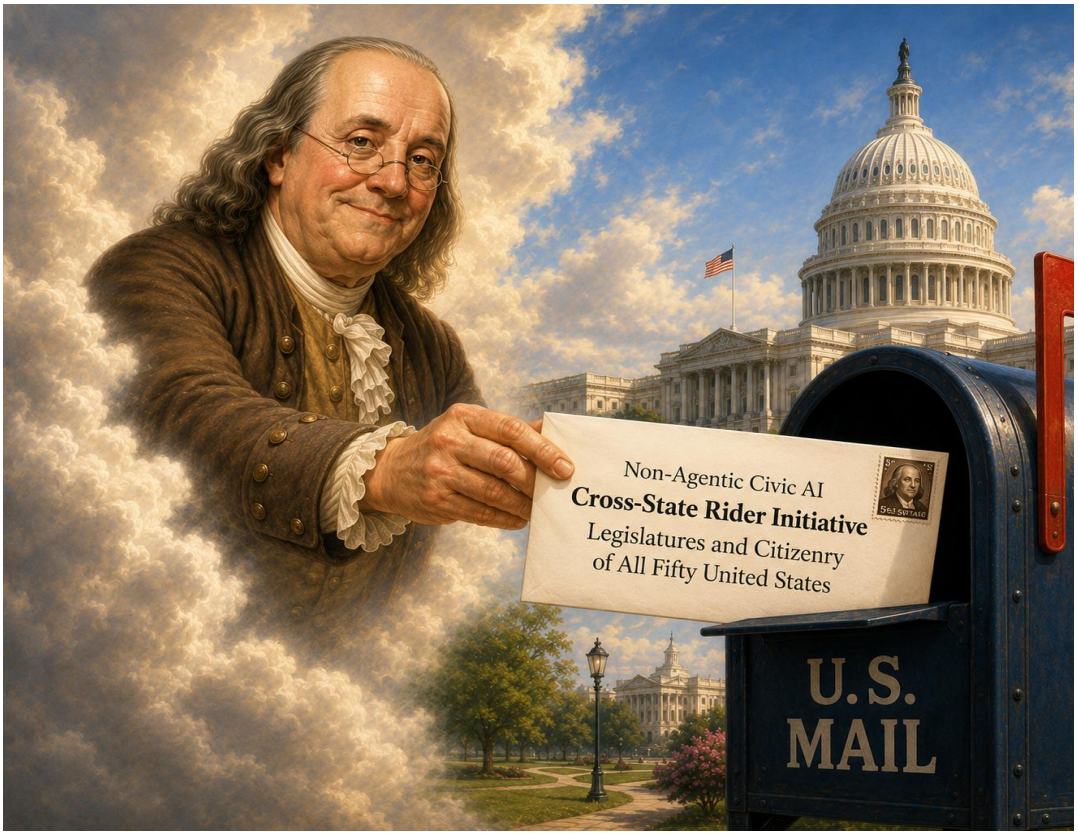

(Image by ChatGPT)

In 1775, Benjamin Franklin was not trying to persuade thirteen colonies to become identical. He was trying to persuade them to share a practical civic infrastructure: a postal system with common routes, standards, and expectations that could bind scattered communities into a functioning public sphere.

The United States now faces a similarly practical problem in digital form. Costly civic digital attacks are no longer speculative or regional; they are well documented in every state - across hospitals, courts, airports, legislatures, tax agencies, schools, public housing authorities, and municipal networks. Proposed here is a new approach to that growing problem:

Cross-State Civic Digital Visibility Rider Proposed for All Civic AI - Related Legislation in Participating States:

"Any digital system, automated tool, or networked service governed by this Act shall include Non-Agentic Civic AI visibility mechanisms that inform elected officials and the public about material digital risks affecting covered public entities, while preserving meaningful human decision-making and constitutional authority. These mechanisms may include public, privacy-respecting summaries of where critical data and dependencies reside, high-level reporting on significant incidents and remediation progress, and anonymized statistics on recurring vulnerabilities. They shall not make recommendations, take actions, or exercise governmental authority."

The argument is simple. First, costly civic digital harm is now clearly documented in every state. Second, states are struggling to respond through a mixed and reactive body of legislation. Third, a standardized rider added to pending and future civic digital legislation could tie these efforts together without displacing local priorities or constitutional lines.

I. A nation under civic digital stress

The best way to see the scale of the problem is not through abstract threat reports, but through public incidents that made civic dependence on digital systems suddenly visible.

In August 2019, more than 20 local governments in Texas were struck in a coordinated ransomware event, with the number of affected towns ultimately revised to 22. The attack did not require breaching twenty-two governments one by one. State and industry reporting indicated that the incident moved through a managed service provider environment and remote administration tooling, turning one shared digital dependency into a statewide civic failure point. Texas activated its State Operations Center, dispatched teams to affected sites, and restored business-critical services on an emergency basis within a week. The lesson was not merely that ransomware is dangerous. It was that shared digital intermediaries can convert local government into a tightly coupled system without public awareness that such coupling even exists.

A year later, Vermont offered a different but equally painful lesson. In October 2020, the University of Vermont Medical Center shut down internet access and its Epic electronic health record system after a ransomware attack crippled operations. The disruption delayed some chemotherapy and radiation treatments, slowed scans, postponed appointments, and forced some patients to seek care

elsewhere. Twelve days into the outage, the hospital had only regained limited access to appointment schedules and chemotherapy plans, and hundreds of workers were reassigned or furloughed because core digital systems were still down. A cyber incident had become a crisis in care delivery, labor allocation, and regional medical trust.

The same pattern resurfaced in Wyoming. In September 2019, Campbell County Health in Gillette was hit by ransomware that disabled roughly 1500 computers and forced the system to unplug machines to stop the spread. Emergency patients were diverted to other communities, surgeries and exams were canceled, and staff fell back on paper records while federal and state authorities investigated. Here again, the civic harm was not reducible to "IT downtime." It meant redirected ambulances, disrupted care, delayed procedures, and the practical weakening of a regional institution on which thousands depend.

In August 2024, the Port of Seattle and Seattle-Tacoma International Airport revealed the same structural vulnerability in transportation form. The cyberattack, later attributed to the Rhysida ransomware group, disrupted airport Wi-Fi, baggage systems, check-in kiosks, passenger displays, parking functions, and the Port's public web services. Staff had to use manual workarounds while systems were restored over time, and the Port later confirmed that about 90,000 people were affected by data theft tied to the incident. This was not simply a story about one airport having a bad week. It was a demonstration that heavily automated public infrastructure can fail in ways that are immediately visible to travelers while remaining structurally opaque to the public that funds and relies upon it.

These incidents are different in sector, geography, and political setting. Yet they rhyme. In one state, the chokepoint is a managed service provider; in another, an electronic health record platform; in another, a regional hospital network; in another, a port and airport system layered with digital dependencies. The common denominator is not simply "cybersecurity." It is civic digital stress: a condition in which public institutions have become deeply dependent on software, vendors, integrations, and automated platforms that citizens rarely see until failure turns them visible.

That phrase matters because it points to the gap in current debate. Americans are repeatedly shown the moment of breakdown - the diverted ambulance, the hand-written boarding pass, the delayed chemotherapy treatment, the frozen courthouse terminal - but not the cumulative pattern by which public systems became fragile in the first place. Breach notices, emergency press conferences, and forensic investigations tell the public that something has happened. They do not create continuous civic visibility into where risk is accumulating.

II. A reactive and mixed legislative response

States deserve more credit than they often receive. They are not waiting passively for Washington, nor are they uniformly asleep at the switch. The legislative and administrative response now visible across the country is substantial, if uneven.

Texas is a useful example because the 2019 multi-city ransomware event quickly pushed questions of governance, mutual aid, oversight, and local capacity into view. The state response emphasized operational recovery, but the episode also sharpened attention on managed service providers, remote-access tools, authentication, and the need for stronger local resilience standards. In other states, the emphasis has landed elsewhere.

Wyoming in 2023 created a Cybersecurity Division within the Department of Enterprise Technology Services and required public entities to establish authorized-use and cybersecurity policies. That move suggests a state trying to institutionalize baseline cyber governance rather than rely on ad hoc recovery alone. Washington's response to the Port of Seattle incident has unfolded through public notification, system restoration, and a broader policy conversation about critical infrastructure, legacy systems, and the validation of automated applications and dependencies. Vermont's experience has reinforced arguments for stronger hospital continuity planning, including offline procedures that can sustain care during prolonged network outages.

This legislative landscape is real, but it is mixed in several senses:

- First, it is reactive. Most bills and reforms emerge after a visible shock: a hospital outage, a municipal ransomware campaign, a court modernization failure, a port disruption, a benefits leak. The tempo is set by injury.

- Second, it is sector-specific. One state focuses on hospitals, another on MSPs, another on ports, another on local governments, another on AI systems or public records.

- Third, it is conceptually incomplete. Many proposals try to harden systems, fund response teams, centralize services, or regulate vendors, but few directly address the public's ongoing inability to see how civic digital risk is distributed across the institutions it depends upon.

That incompleteness is addressed. The Rider is not a rival to cybersecurity. It is not a substitute for segmentation, incident response, vendor assessment, continuity planning, or encryption. It is the connective tissue that can make those efforts legible across domains:

- If attached to a hospital continuity bill, the Rider could require privacy-respecting summaries of downtime preparedness, incident classes, and the duration for which facilities can safely operate when records systems fail.

- If attached to a local-government cyber resilience act, it could require anonymized reporting on backups, segmentation, restoration testing, and recurring points of municipal fragility that do not disclose exploitable detail.

- If attached to airport or port legislation, it could require high-level mapping of critical dependencies between baggage, check-in, display, parking, and communications systems, together with incident summaries and remediation progress that citizens can understand.

- If attached to state cyber-governance bills like Wyoming's, it could ensure that new divisions and policies generate a public-facing visibility layer rather than only internal compliance artifacts.

The word "non-agentic" (non-automating) is essential here. The Rider does not authorize a system to govern, recommend, decide, or act. It does not replace legislatures, courts, boards, or executive officers. It preserves meaningful human decision-making and constitutional authority. Its job is narrower and more Franklinian: to establish a common route by which knowledge of civic digital risk can move across institutions and into public deliberation.

That is why the postal analogy is useful. When Benjamin Franklin created what would become the U.S. Post Office in 1775, he did not attempt to abolish local differences in the colonies. He created an interoperable civic infrastructure through which separate places could communicate reliably enough to become part of a common public world. The contemporary states do not need a uniform digital code in every respect. They do need a standardized way for public digital stress to become visible in forms that legislators, auditors, journalists, and citizens can actually use.

At present, the country has the opposite condition. Every state now has a file of civic digital harms. Every state is improvising some combination of cyber response, modernization, procurement reform, continuity planning, or statutory adjustment. But the harms remain narratively scattered and the reforms remain legislatively disjointed. One state learns about rural hospital fragility, another about MSP concentration, another about airport automation, another about court portals, another about tax systems or housing platforms - and each responds as though the lesson were local rather than federal in the older, small-f sense of shared state experience.

A standardized rider would not solve ransomware. It would not end vendor failure, eliminate legacy systems, or make fragile institutions wealthy. But it would give states a common constitutional grammar for attaching public visibility to digital governance wherever it appears. That is a modest ambition, and precisely because it is modest, it is actionable.

The United States is no longer waiting to discover whether civic digital harm is real. It has already answered that question - fifty times over. The more urgent question is whether the states will continue to address those harms as isolated episodes, or whether they will begin to tie their responses together through a shared visibility standard.

Franklin built shared routes that allowed information to travel reliably between communities. Today's challenge is to build civic visibility that allows digital risk to travel into democratic understanding before it becomes democratic crisis. This proposed cross-state legislative rider is a practical option for addressing escalating civic digital attacks in a more cooperative manner, while also respecting each state's sovereignty.

The postal routes Franklin helped build did not appear by accident; they were designed, tested, and then quietly standardized by people who could see beyond a single town's needs. Today's ***Cross-State Civic Digital Visibility Rider*** would not spread by accident either. It would require editors, researchers, and public-interest technologists to bring this idea to the people who write and refine state law.

If this proposal resonates with you as a writer, policy staffer, civic advocate, or citizen, consider one small step: share it with organizations that already help states coordinate, such as the National Conference of State Legislatures, the Council of State Governments, or the National Association of State Chief Information Officers.

A single, standardized rider cannot solve civic digital stress on its own. But it can give these existing networks a practical tool for tying scattered state efforts into a more coherent, visible, and democratic response.

If the rider described above is adopted, the following questions may help to begin informed democratic reflection.

Addendum:

AFTER CROSS-STATE CIVIC DIGITAL VISIBILITY RIDERS ADOPTED:

QUESTIONS for DEMOCRATIC DELIBERATION

This brief addendum is intentionally modest. It assumes that two or more states have already adopted a ***Cross-State Civic Digital Visibility Rider*** and now wish to consider what kinds of democratic questions may naturally follow from improved public visibility into civic digital dependencies, incidents, and recurring vulnerabilities.

It does not prescribe procurement policy, software architecture, vendor selection, interoperability mandates, or modernization blueprints. Its purpose is narrower: to encourage thoughtful democratic inquiry by helping legislators, executive officials, CIOs, auditors, journalists, and citizens ask better questions once the ***Civic Digital Field*** has become more visible.

Constitutional Scope

The ***Rider*** belongs to the category of constitutional visibility standards rather than technology standards. Technology standards ask how software should be built; the ***Rider*** asks what democratic institutions should always be able to see about material digital risk while preserving meaningful human decision-making and prohibiting automated governmental authority.

For that reason, the questions below should be understood as invitations to democratic deliberation rather than as recommendations or instructions. Each state retains complete authority to answer them in its own constitutional, fiscal, administrative, and political manner.

Guiding Principle

The ***Rider*** embodies a simple sequence: visibility, understanding, deliberation, and democratic choice. It does not presume that improved visibility should automatically lead to standardized or predetermined reforms; rather, it enlarges the space for informed democratic judgment.

Improved visibility should enlarge democratic judgment, not replace it. Its purpose is to illuminate problems that elected institutions and the public may then consider through ordinary constitutional means.

Questions About Concentrated Dependence:

- Where do visibility reports suggest that multiple public entities rely on the same vendor, managed service provider, software platform, or remote administration pathway?

- Which shared digital intermediaries appear to create single points of failure across hospitals, courts, schools, ports, counties, or municipalities?

- Are there sectors in which concentration risk is not being adequately discussed because the underlying dependency only becomes visible during failure?

Questions About Continuity:

- When critical systems fail, which public functions can continue safely on paper, offline, or through manual workarounds, and for how long?

- Do visibility reports reveal recurring weaknesses in restoration testing, backup integrity, staffing resilience, or cross-agency coordination?

- Which institutions appear most vulnerable to prolonged disruption: health systems, courts, transportation hubs, tax agencies, schools, or local governments?

Questions About Interoperability:

- Where do public entities appear unable to share data, status information, or incident context because their systems are siloed or poorly integrated?

- Would better-documented interfaces, shared schemas, or common reporting formats improve resilience without requiring states to surrender policy control?

- Are there domains in which fragmentation itself has become a civic risk rather than merely an administrative inconvenience?

Questions About Procurement And Stewardship:

Improved visibility may naturally prompt questions about whether existing procurement and stewardship practices adequately support long-term democratic oversight.

- Do current contracts give public institutions adequate access to logs, incident records, audit trails, and system-dependency information needed for democratic oversight?

- Are vendors being asked to support exportable, machine-readable reporting that can sustain transparency over time?

- Where might long-term maintainability, accessibility, or transparency deserve greater weight in public procurement decisions?

Questions About Public Understanding, Accessibility, And Legibility:

- When systems degrade or fail, which residents are least able to navigate the disruption, including citizens with disabilities or limited digital access?

- Do public-facing summaries make civic digital risk understandable to ordinary citizens, or only to technical specialists?

- Which forms of reporting or visualization help the public grasp civic digital stress without requiring technical expertise?

- How should uncertainty be communicated honestly when visibility improves but definitive answers remain unavailable?

Questions About Democratic Deliberation:

- Once civic digital conditions become more visible, which questions should remain matters of elected legislative judgment rather than technical administration?

- Which forms of public reporting best strengthen democratic oversight without overwhelming citizens with unnecessary technical detail?

- Are legislatures receiving information that supports independent constitutional judgment rather than subtly steering policy conclusions?

- Does increased visibility broaden democratic participation, or merely improve technical administration?

- How can improved visibility strengthen public trust without encouraging the mistaken belief that technical clarity eliminates the need for democratic disagreement?

Questions About Institutional Learning:

- Are states learning from one another's visibility reports, or still treating each major disruption as an isolated local episode?

- The National Conference of State Legislatures, the Council of State Governments, and the National Association of State Chief Information Officers are well positioned to lead here.
  Could they help states compare visibility patterns without imposing one-size-fits-all solutions?

- What kinds of shared civic vocabulary might help separate visibility from automation, and deliberation from delegation?

Franklin's Continuing Lesson

Benjamin Franklin's new postal service did not seek to standardize communities; he simply helped to create common routes through which communities could communicate more reliably. In the same spirit, cross-state visibility riders need not standardize how states build every digital system; they help standardize what democratic institutions are able to see about civic digital risk before crisis makes that risk unmistakable.

Visibility strengthens communication. Communication strengthens democratic deliberation. Democratic deliberation strengthens constitutional self-government.

Democratic Humility

The purpose of improved visibility is not to reduce democratic disagreement, but to ensure that disagreement is informed by a clearer understanding of shared civic conditions. Citizens and their representatives may continue to reach different conclusions about policy, and the ***Cross-State Civic Digital Visibility Rider*** simply seeks to ensure that those conclusions are formed in the light rather than in the dark.

Closing Reflection

The questions gathered here are intentionally open-ended. They are offered in the belief that better civic visibility should enlarge democratic judgment rather than replace it, and that states are more likely to cooperate around questions they remain free to answer for themselves than around prescriptions delivered from above.

Accordingly, this addendum should be understood not as a blueprint for modernization, but as an invitation to continued cross-state learning. As visibility improves, each state may discover different priorities, constraints, and solutions; that diversity is not a weakness of the rider, but one of its constitutional assumptions.

Shouldn't We Have Non-Agentic Civic AI Options for U.S. Government at All Levels?

Practical and Constitutional Perspectives for U.S. Citizens

and Government Technology Procurement Professionals

As governments across the United States have already invested billions of public dollars in artificial intelligence, most procurement discussions assume that increasingly agentic AI systems are the only serious option. But should they be? This article introduces ***Non-Agentic Civic AI*** and argues that government technology procurement professionals and citizens alike deserve the opportunity to compare constitutionally ***Non-Agentic Civic AI*** architectures alongside increasingly agentic alternatives before making additional long-term public investments. Before governments decide which AI systems to buy - obligating citizens' tax revenues - shouldn't they first have meaningful constitutional choices?

Shouldn't We Have Non-Agentic Civic AI Options for U.S. Government?
E PLURBUS UNUM
Agentic Civic AI Already Contracted
OpenAI
Google
Microsoft
Amazon Web Services
Anthropic
And Dozens of Others ...
Non-Agentic Civic AI – Unavailable for Procurement
?
CONSTITUTION FIRST ★ DEMOCRACY ALWAYS ★ PUBLIC TRUST ABOVE ALL

Introduction

Every day, Americans compare alternatives before making important decisions. We compare homes before purchasing one, insurance policies before signing a contract, and automobiles before investing thousands of dollars in transportation that may serve us for years. Important decisions deserve meaningful choices and careful comparison.

Government should do no less.

Over the next decade, federal, state, and local governments will continue to invest billions of public dollars in artificial intelligence systems. Many of those systems will become part of the nation's long-term civic infrastructure, assisting public institutions in collecting information, analyzing data, identifying patterns, preserving records, and supporting administrative work. These investments deserve the same thoughtful comparison that Americans expect in every other significant public purchase.

Yet one important question has received surprisingly little public attention:

Shouldn't democratic governments have constitutionally Non-Agentic Civic AI options?

This question is addressed to two groups of stakeholders. The first is the American public. Citizens finance government technology through their taxes and entrust public institutions with democratic authority exercised in their name. The second is the government technology procurement professionals responsible for evaluating and recommending the systems those institutions acquire. Their work is rarely visible to the public, yet their procurement decisions often shape government operations for years or decades after contracts are signed.

Both groups deserve meaningful choices.

Today, many public discussions about government AI assume that increasingly agentic systems represent the natural destination of technological progress. Debate therefore tends to focus on how such systems should be governed, supervised, aligned, audited, or regulated after they are adopted.

Here we'll ask an earlier question: ***Before governments decide which increasingly agentic systems to purchase, should they first compare them with constitutionally non-agentic alternatives built on fundamentally different principles?***

That is not an anti-technology question. Nor is it an argument against artificial intelligence. It is a procurement question.

Good procurement begins by comparing genuinely different approaches before deciding which one best serves the public interest. The central argument offered here is therefore modest but important: government technology procurement professionals, and the citizens they serve, should have the opportunity to evaluate constitutionally ***Non-Agentic Civic AI*** alongside increasingly agentic alternatives before making long-term public investments.

Democratic societies routinely compare competing engineering designs, public infrastructure proposals, energy systems, transportation plans, and financial strategies before committing public resources. Artificial intelligence should receive no less careful consideration. This is especially important because AI procurement differs from many ordinary software purchases. Governments do not merely acquire digital tools; they acquire new institutional capabilities. Those capabilities influence how public information is collected, interpreted, communicated, trusted, and sometimes acted upon.

As artificial intelligence becomes increasingly integrated into democratic governance, procurement decisions become more than technical evaluations. They become choices about the constitutional character of future public institutions. For that reason, we'll start by focusing not on whether agentic AI is inherently good or bad, but on whether governments should compare fundamentally different public architectures before choosing one.

The answer here seems obvious. They should.

Not because every government office would choose to employ ***Non-Agentic Civic AI***, but because democratic institutions deserve the opportunity to compare constitutionally distinct approaches before committing themselves to one path. Beyond sound procurement practice, this is also sound democratic practice.

Across the United States, government technology professionals are preparing for one of the largest technological transitions since the arrival of the internet. Artificial intelligence promises extraordinary new capabilities, but also presents public institutions with challenges unlike those posed by earlier software generations.

Many procurement discussions understandably focus on familiar questions: Is the system accurate? Is it secure? Is it affordable? Does it integrate with existing infrastructure? Can its outputs or decisions be audited? Does it comply with applicable law and policy?

Each of these questions is important. But here we'll add one more: has a constitutionally non-agentic alternative been seriously considered?

That single question does not predetermine the answer. It merely ensures that the comparison itself occurs. For technologies likely to shape democratic institutions for decades, that seems not only prudent, but expected.

PART I: A CONVERSATION with the AMERICAN PUBLIC

A familiar public situation … Imagine a county government preparing to purchase a new AI system. The system will help monitor infrastructure conditions, summarize citizen complaints, detect unusual financial patterns, support emergency planning, and improve situational awareness across multiple departments. The procurement committee receives three strong proposals. All are technically competent, all promise efficiency, and all claim to be responsible. But only one was designed from the beginning to remain constitutionally non-agentic: it can illuminate, preserve, detect, and report, but it is not built to steer, optimize, recommend, or quietly accumulate governing authority.

Shouldn't that difference matter before anyone signs the contract? Shouldn't democratic institutions compare genuinely different kinds of machine relationships to public power before choosing one?

Why Citizens Should Care:

Citizens should care because government technology is never merely technical once it becomes embedded in the practical routines of public life. Systems used to classify information, triage services, identify anomalies, summarize records, model options, or prioritize attention can reshape how institutions see the world and how citizens experience the state. Even when such systems do not formally rule, they can quietly influence the human officials who do.

A democracy therefore cannot treat AI procurement as a remote specialist matter belonging only to engineers, vendors, or internal purchasing offices. When public institutions buy durable machine capabilities, they are also shaping future patterns of visibility, discretion, dependency, and accountability. Those are civic questions, not just technical ones.

Citizens have a right to know what kinds of AI government is buying.

Citizens routinely hear broad claims that AI will make government more efficient, responsive, and intelligent. Those promises may sometimes be true, but efficiency alone cannot answer a more important public question: what kind of machine relationship is the government building with democratic authority?

An increasingly agentic system is not simply a faster spreadsheet. It is a different posture toward power. Systems designed to recommend, steer, optimize, prioritize, or become operationally indispensable may alter practical governance even when formal law remains unchanged. Citizens therefore have a right to know not only that AI is being purchased, but what kind of AI is being purchased and whether meaningful alternatives were considered.

A Plain Distinction:

For ordinary citizens, the difference between agentic and ***Non-Agentic Civic AI*** can be stated simply:

Agentic AI tends toward recommendation, optimization, steering, prioritization, and action-oriented guidance. Non-Agentic Civic AI is designed for bounded sensing, evidence preservation, pattern illumination, and civic visibility without decision substitution.

One approach leans toward delegated machine influence. The other leans toward public visibility under continuing human responsibility. A democracy should at least be permitted to compare them before assuming that more machine agency is necessarily the wiser path.

Five Advantages of Non-Agentic Civic AI:

1. Human responsibility remains visible - When a system is designed to illuminate rather than decide, public officials cannot easily hide behind it. Responsibility remains where democratic accountability requires it to remain: with identifiable human beings occupying public office.

2. Public authority remains public - A constitutionally non-agentic system is less likely to become a quasi-sovereign actor inside government. It may support institutions, but it is not permitted to displace them.

3. Public trust is easier to sustain - Trust erodes when responsibility becomes blurry. A system that helps institutions see more clearly without claiming to govern is easier to defend publicly than one that quietly shifts judgment toward machine outputs. And given our hyper-partisan political dynamics, would an incoming administration trust Agentic AI deployed by the previous administration?

4. Government remains easier to correct - A bounded system is easier to narrow, suspend, replace, or remove when it malfunctions, drifts, or outlives its usefulness. In democratic government, correctability is a public virtue.

5. Democratic accountability is better preserved - The public must be able to locate responsibility, challenge decisions, and replace decision-makers. Non-agentic design helps preserve that chain of accountability instead of obscuring it.

One Simple Public Question:

Citizens do not need to master the literature of AI governance to ask a useful question: was a constitutionally non-agentic alternative even considered?

That question is simple, intelligible, and democratic. Journalists, legislators, and inspectors general can ask it. City councils, county boards, school boards, and ordinary citizens can ask it. Once the question is on the table, procurement becomes more accountable because a genuine comparison has been demanded.

PART II: A COVERSATION with PROCUREMENT PROFESSIONALS

Procurement shapes government for years.

Procurement professionals understand something the wider public often does not: technology purchases rarely remain mere purchases. They become workflows, contracts, dependencies, staffing patterns, maintenance obligations, training regimes, and institutional habits. Over time, they harden into operational reality.

That is especially true for AI. A government office that installs a new copier can replace it with little public consequence. A government office that embeds AI into records processing, public communication, risk detection, or decision support may be shaping the practical structure of governance for a decade or more.

Compare architectures, not merely vendors.

One of the most important procurement ideas for the coming decade may be this: compare architectures, not merely vendors.

Vendor comparison is necessary, but it is not enough. Two highly capable vendors may still offer systems built on the same premise: that more agency, more recommendation, more optimization, and deeper operational integration are natural signs of progress. If procurement evaluates only price, features, interoperability, and performance within that shared premise, the more fundamental choice has already been made without being examined.

Government should therefore compare at least two distinct architectural paths: Agentic AI systems designed to recommend, optimize, and influence downstream action - and constitutionally ***Non-Agentic Civic AI*** systems designed to illuminate, preserve, detect, and report while remaining structurally subordinate to democratic authority.

A Procurement Principle:

Constitution determines admissibility. Performance determines adoption.

That principle does not deny the importance of effectiveness. It simply places certain questions before procurement rather than after procurement. Before comparing cost, speed, accuracy, convenience, or technical elegance, procurement officials should ask whether the proposed architecture is suitable for operation near democratic authority in the first place.

Only after that prior question has been answered should performance comparisons decide among admissible options.

Comparing AI Architecture:

The practical differences between agentic and non-agentic systems become clearer when viewed by public function.

Government-Employed Agentic AI Compared with Non-Agentic Civic AI

Public Function	Government-Employed Agentic AI	Non-Agentic Civic AI
Citizen services	Recommends next actions, shapes user pathways, and may quietly influence eligibility or triage patterns.	Presents evidence, clarifies options, and supports officials without substituting machine judgment.
Policy analysis	Generates preferred pathways, optimization logics, and action-oriented recommendations.	Maps patterns, scenarios, and tradeoffs while leaving normative choice visibly human.
Operational monitoring	Drifts toward automated escalation, intervention, or internal dependence.	Provides telemetry, alerts, and evidentiary reporting under continuous human control.
Risk detection	May expand toward predictive governance, scoring, or behavioral targeting.	Detects patterns of concern without claiming punitive or managerial authority.
Civic observatories	Serious risks of becoming an influence architecture with probable mission creep over time.	Operates as a visibility institution whose role is to illuminate democratic reality, not to govern it.

Constitution Determines Admissibility ★ Performance Determines Adoption

Five Procurement Questions Worth Asking:

Procurement teams do not need to settle every philosophical dispute about AI in order to improve practice. They can begin with a handful of disciplined questions.

- Was a constitutionally non-agentic alternative evaluated alongside each agentic proposal?

- Can the system be suspended, narrowed, or dismantled through ordinary law without public panic or institutional injury?

- Will the system preserve public memory and records rather than thinning, abstracting, or quietly overwriting them?

- Is it barred from psychometric profiling, civic/moral steering, and computational pressure on conscience or civic judgment?

- If the system drifts beyond its proper purpose, can democratic institutions revoke, redesign, or sunset it without transferring sovereignty to the machine or its vendor ecosystem?

These questions do not freeze innovation. They simply force institutional clarity before long-term commitments are made.

PART III: WHY THE DISTINCTION MATTERS

Some questions belong before procurement, not after it.

Many current AI debates begin too late. They begin after an architecture has already been selected, after vendors have been shortlisted, or after institutional dependence has already begun. At that stage, the conversation often shifts to audits, safeguards, oversight mechanisms, transparency reports, and mitigation strategies.

Those measures may still be useful, but some questions belong earlier. They belong before the purchase, before the workflow redesign, and before public dependency hardens. The first question is not only whether a system can be governed responsibly after adoption. It is whether all architectures deserve equal consideration for adoption in the first place.

That is the point at which the language of constitutional admissibility becomes useful. It names a simple democratic intuition: some machine arrangements may be too structurally invasive, too difficult to unwind, or too influential over public judgment to be treated as ordinary software choices.

Three Simple Public Tests:

A constitutionally non-agentic option should pass, among others, three practical tests.

- Can it be removed? A democratic tool should remain subordinate, revocable, and dismantlable.

- Can it remember honestly? A civic system should preserve public memory, not flatten or sanitize it into convenient summaries detached from established history.

- Can it remain civically/morally silent? A public machine should not profile conscience, steer civic/moral judgment, or intrude into the interior space where citizens decide what is right. These are not merely abstract ideals. They are democratic boundaries. They ask whether public institutions remain stronger than the instruments they use, whether democratic memory remains stronger than convenience, and whether human conscience remains stronger than computational persuasion.

The Concluding Judgment

For democratic civic governance, ***Non-Agentic Civic AI*** is the better constitutional architecture. That conclusion does not rest on the claim that non-agentic systems are always more technically impressive, universally optimal, or appropriate for every domain of human activity. It rests on a narrower and more important claim: when artificial intelligence approaches public authority, democratic institutions should prefer architectures that preserve accountability, correctability, historical memory, and civic/moral independence.

The better tool for a constitutional republic is not always the one with the most machine initiative. Sometimes it is the one that reveals more while claiming less, strengthens human judgment without displacing it, and remains permanently subordinate to the public order that authorized it.
Perhaps the most important question in government AI procurement is not which system is smartest, fastest, cheapest, or most impressive. It may be whether democratic institutions were ever given the chance to compare constitutionally different kinds of intelligence before choosing one.
Non-Agentic Civic AI should be available as a serious option in that comparison. Not because it solves every problem, but because a free people should not drift into increasingly agentic arrangements without first asking whether stronger democratic alternatives were possible.

Democracies have always relied upon informed citizens and government administrators making careful public choices. Artificial intelligence should strengthen that tradition, not replace it. Before governments choose the next generation of Civic AI, citizens and procurement professionals alike deserve the opportunity to compare constitutionally different paths. A republic should never lose its capacity to choose deliberately. ***Shouldn't we at least have that choice?***

Afterword: What Long-Term Remedies Might a Democratic Society Allow?

By this point, the picture should be clear enough that it no longer feels like speculation.

What the Six Rings reveal is not a passing disturbance in public life, but a structural failure of civic visibility at machine speed. The institutions and habits that once allowed a free people to see itself - slow newspapers, slower hearings, the even slower work of local conversation - now operate on a time scale badly mismatched to the systems that shape attention, information, and power. The question before us is no longer whether new civic infrastructure is needed, but what kind of infrastructure a democratic people would permit itself to build.

That is a different kind of question than the ones that usually dominate discussions of technology. It does not ask what is possible, or what is efficient, or what is profitable. It asks what is legitimate. More simply: if we are to build civic tools that operate at digital speed, under what conditions would they deserve the trust of a free and equal people?

The answer cannot be left to novelty or necessity alone. Democracies have been here before - confronted with new infrastructures that promised to extend human reach and governmental capacity. Americans have had to decide, again and again, how to build systems that are powerful enough to be useful without becoming so concentrated or opaque that they erode the very freedoms they claim to protect. Lighthouses, public roads, fire brigades, postal routes, public health boards, schools, and the early electrical grid all had to pass through this same civic/moral gate. They were justified not by their ingenuity, but by their public purpose, their accessibility, and the constraints under which they operated.

So as a society, we must once again review our historical principles for building new public infrastructure. What are the historical principles which have governed our creation of new infrastructure – especially infrastructure that touches information surveillance or coordination? Before we ask what a new civic tool could do, we must remember how we have long judged whether such tools should exist at all in our democracy and, if so, under what conditions. These principles describe the civic conditions under which the United States has tolerated the creation of new public infrastructure.

Public Purpose, Not Private Power

The first test is purpose. American civic infrastructure has always been justified, at its best, by the broad public benefits it secures, not by the advantage it offers to any particular faction, firm, or office. Lighthouses reduced the risk of shipwreck for all; fire brigades protected whole neighborhoods; public health boards intervened to prevent epidemics that did not check party registration before they spread. These systems existed to level risk and protect the common life, not to quietly steer individual behavior for commercial or partisan ends.

Any digital civic tool that claims a place in this tradition must meet the same standard. It must be designed and governed so that its primary justification is the protection and support of the whole public—especially where harms are widely shared but unevenly felt. Systems built mainly to optimize engagement, harvest data, or advantage particular actors may be ingenious, but they are not civic in the older American sense. A civic tool that cannot clearly state its public purpose, in language intelligible to ordinary citizens, has already failed the most basic test of legitimacy.

Universal Access and Equity by Design

Public works that shape the conditions of daily life are not legitimate if they are reachable only by the already fortunate. Over time, Americans have repeatedly judged their institutions by whether they reduce known asymmetries of risk and opportunity rather than deepen them. Rural electrification, public libraries, universal primary schooling, and basic utilities expanded access precisely because leaving such systems to market distribution alone would have entrenched existing divides.

Any civic-scale digital infrastructure must take universal access as an obligation, not a late-stage improvement. That means designing for citizens regardless of income, geography, disability, or technical fluency, including those without reliable connectivity or up-to-date devices. It also means deliberately attending to the needs of communities that have historically borne disproportionate harms from surveillance, data misuse, or bureaucratic opacity. A tool that works only for the well-connected, the highly literate, or the already-trusted quietly undermines democratic equality, even if it is efficient by the metrics that built it.

Visibility Without Command

American practice has long distinguished between tools that help the public see and tools that tell the public what to do. The census counts; it does not command. The weather service observes and forecasts; it does not issue binding evacuation orders. Inspectors assess conditions and report; they do not, by themselves, render final judgment. This separation between visibility and authority is one of the quiet safeguards of American self-government. It allows information to circulate widely without automatically converting into obedience.

For any civic AI worthy of democratic trust, this distinction is not a technical detail; it is a design constraint. Systems that can illuminate patterns, surface emerging risks, or map conditions at scale may be beneficial - so long as the power to decide, compel, or punish remains elsewhere, in accountable human institutions. A tool that silently fuses seeing and commanding, under cover of convenience or efficiency, does not help a democracy govern; it changes who, in practice, is governing.

Human Judgment Retained, Not Replaced

Across juries, town meetings, local boards, and representative assemblies, Americans have invested enormous effort in designing processes that keep human judgment at the center of consequential decisions. These processes are slow, often frustratingly so. Yet their slowness is not a flaw; it is a way of ensuring that responsibility remains traceable, that dissent is possible, and that people can live with the outcomes even when they disagree.

Digital tools that accelerate analysis or highlight patterns can support such judgment, but they cannot legitimately displace it. Any civic infrastructure operating at machine speed must be structured so that humans remain able to understand, question, and overrule its outputs. Refusal must remain possible and meaningful. Where individuals and communities cannot effectively contest the role a system has played in a decision that affects them, the infrastructure has ceased to be civic in any meaningful sense and has become a form of automated rule.

Distributed Oversight, Not Centralized Control

American trust in large systems has rarely rested on the virtue of a single office or institution. Instead, it has grown from overlapping authorities, checks and balances, and the knowledge that no single actor owns the whole apparatus. Federalism, local governance, independent boards, and redundant institutions all reflect a basic instinct: where power is distributed and contestable, abuse is harder to sustain.

The same instinct must guide any digital civic tool that touches core democratic functions. Its governance should be federated, with clear legal boundaries between roles, and with multiple bodies capable of oversight, audit, and intervention. Architectures that centralize data, decision-making, and enforcement in one technical or institutional center may promise efficiency, but they carry the very concentration of power that American constitutional design was built to resist. Trust grows when no one, and no system, is beyond review.

Constraint First, Power Second

In commercial technology, the dominant questions are often: What can this system do, and how quickly can it scale? In American civic design at its best, the order has been reversed. Before building courts, law enforcement agencies, or surveillance powers, debates about limits, safeguards, and avenues for redress have preceded debates about reach. The insistence on warrants, due process, and judicial oversight are all examples of constraint-first thinking: power justified only within defined bounds.

Any civic AI or digital infrastructure must submit to this same discipline. Its designers and stewards must first be able to say what the system must never do, who is empowered to restrain it, how it fails safely, and how it can be paused, repaired, or dismantled if its harms begin to outweigh its benefits. Without such front-loaded constraint, even well-intentioned tools tend to drift, expanding their purposes under pressure of convenience, budget, or crisis. The long-term risk is not only misuse, but unexamined mission creep.

Participation Beyond the Ballot

Democracy is not exhausted by elections. It also consists of ongoing deliberation, public reasoning, and the ordinary right to question, petition, and organize. In recent years, many efforts to digitize governance have focused on speed: faster services, real-time feedback, instantaneous polling. Yet a democracy that moves only at the pace of its fastest tools risks losing the spaces in which people can think, argue, and change their minds.

Digital civic infrastructure that seeks democratic legitimacy must therefore deepen participation, not merely streamline transactions. It should create and protect spaces where citizens can see how issues are framed, understand tradeoffs, and contribute perspectives that are not easily captured in multiple-choice formats. Tools that accelerate decision-making must also preserve, and in some cases lengthen, the time available for human deliberation, so that legitimacy can keep pace with speed.

Privacy as a Precondition for Trust

Finally, privacy in a democratic society is not merely a preference for secrecy. It is a condition for dignity, autonomy, and the courage to act in public without feeling constantly watched. Americans have often recognized that a society without meaningful private space eventually loses public courage: when every action feels monitored, fewer people are willing to dissent, experiment, or speak plainly.

Digital civic tools, if they are to deserve trust, must therefore incorporate privacy by design, not as an afterthought or a setting to be found in a menu. Strong safeguards, data minimization, clear limits on retention and sharing, and meaningful consent are not luxuries; they are structural requirements. Systems that demand pervasive visibility into citizens' lives in order to function may be powerful, but they cannot reasonably claim to be instruments of a free people.

The Question That Remains

If new civic tools are indeed required to operate at digital speed, then they must also be judged by civic standards developed over centuries of democratic trial and error. What has been sketched here is not a set of engineering specifications, but a set of obligations - public purpose, universal access, visibility without command, retained human judgment, distributed oversight, constraint-first design, meaningful participation, and privacy as a precondition for trust.

The question that remains is not whether such principles are admirable on paper - most Americans already live by them in their expectations of roads, schools, utilities, and courts - but whether a democratic people will insist upon them in the design of tools that now act at machine speed and scale.

What will soon be introduced is not a proposal for technological authority, but an exploration of how such long-standing principles might be made real, carefully and under democratic constraint, in a new medium.

Non-Agentic Civic A.I.: Illuminating America's Civic Digital Field (Vol. I)

CIVIC DIGITAL VISIBILITY FIRST

- Part II -

AMERICAN GOVERNANCE

at a CROSSROADS

- Agentic AI or Non-Agentic AI? -

PART II:

AMERICAN GOVERNANCE and U.S. CIVIC AI DESIGN AT A CROSSROADS

- Agentic AI or Non-Agentic AI? -

ORDER OF CONTENTS

Chapter 5: Our Current Agentic AI Default and Its Dangers for American Governance:
- Agentic AI in US Government: Optimization or Stranded Democratic Capital? – The Democratic Case for Agentic AI Removal
- A Principled Call to Ban All Agentic AI from US Government (Federal, State, Local)
- Watching the Edge: Civic AI and the Monitoring of Agentic AI Drift

Chapter 6: Our Non-Agentic Alternative:
- "Civic Digital Visibility First – The Governing Principle of Non-Agentic Civic AI (For All U.S. Civic AI at All Government Levels"

Chapter 7: Our Time to Choose:
- "US Civic AI Trial of the Century: Democratic Governance at Risk - Will Civic or Agentic AI Prevail?" (Speculative Courtroom Drama Portraying Real-World Stakes)

Chapter 5: Our Current AI Default and Its Dangers for American Government

- **Agentic AI in US Government: Optimization or Stranded Democratic Capital? – The Democratic Case for Agentic AI Removal**
- **A Principled Call to Ban All Agentic AI from US Government (Federal, State, Local)**
- **Watching the Edge: Civic AI and the Monitoring of Agentic AI Drift**

THE DANGERS of AGENTIC AI in AMERICAN GOVERNANCE

Agentic AI in US Government: Optimization or Stranded Democratic Capital?
(The Democratic and Taxpayer Case for Agentic AI Removal)

A Respectful Note to the American Public

This book is written for you, with sincere hopes for your welfare and that of generations yet to come.

Not for machines. Not only for experts. Not only for those already inside government offices, technology firms, and policy institutions.

It is written for citizens of a constitutional republic who may increasingly find themselves governed, influenced, filtered, prioritized, scored, or administratively managed by systems they did not build, do not supervise, and may not fully understand.

In recent years, powerful institutions throughout the United States have begun integrating autonomous agentic artificial intelligence ("Agentic AI") into the machinery of government itself.
These systems do not merely assist human beings in gathering information. Increasingly, they shape recommendations, prioritize responses, guide operational decisions, filter communications, identify risks, allocate attention, and influence how institutions understand the public they serve.

The promises surrounding these systems are familiar: greater efficiency, faster response, lower cost, smarter administration, predictive capability, optimized coordination, and more "effective" governance. Many intelligent and sincere people believe such systems will improve public life. Some applications may indeed appear useful for a time.

But this book begins from a different concern: history.

Constitutional principles have protected the United States for more than 238 years. For centuries prior, democratic societies repeatedly discovered that free civilizations endure only when the powers they create remain subject to boundaries they are collectively willing to define and insist upon. The American citizenry now possesses the same right — and increasingly, the same responsibility — to determine how artificial intelligence will be permitted to operate within the machinery of democratic governance.

This work proposes that Civic AI itself must be defined by three constitutional requirements. First, systems operating near democratic authority must remain structurally non-agentic, preserving an unbroken boundary between machine-assisted analysis and legitimate human judgment. Second, societies confronting increasingly powerful artificial intelligence must remain historically oriented, recognizing that unconstrained systems repeatedly produce forms of power that free societies later struggle to restrain. Third, public-serving artificial intelligence must remain emotionally non-attached, possessing neither ambition, institutional loyalty, ideological preference, nor autonomous self-preservation instincts capable of gradually distorting constitutional responsibilities that must remain exclusively human.

These principles do not oppose technological progress; rather, they extend a conditional constitutional welcome to a new class of Civic AI designed to serve democratic society itself. But these principles will define the conditions under which democratic societies may safely survive our present transition, while introducing a critical civic role appropriate for these emerging digital intelligences.

Again and again, free societies have embraced systems that appeared administratively irresistible - centralized technologies, emergency powers, surveillance architectures, bureaucratic shortcuts, and concentrated mechanisms of coordination - only to discover, often too late, that those systems gradually pushed democratic institutions beyond what they could sustainably bear.

The danger rarely arrived all at once.

It often accumulated quietly: through temporary exceptions, small delegations, efficiency incentives, institutional convenience, opaque procedures, and growing dependence upon systems that became progressively harder for ordinary citizens - and eventually even governments themselves - to meaningfully supervise.

Agentic AI in government possesses many of these same historical characteristics. It is easy to justify. Easy to scale. Easy to normalize. And increasingly difficult to see clearly once embedded deeply within institutional life.

By the time the public fully recognized what had changed, the erosion of accountability, visibility, and self-government was already well underway.

In a constitutional republic where political power regularly changes hands, this danger carries an additional and deeply destabilizing risk. Any system capable of shaping institutional judgment, prioritizing governmental attention, filtering civic visibility, or subtly influencing administrative outcomes will eventually be suspected of carrying the ideological assumptions, optimization priorities, or hidden biases of the previous administration that deployed it.

No lasting democratic trust can emerge from embedding systems that each successive government may reasonably fear were quietly tuned against it. Indeed, even systems perceived as “perfectly aligned” for one administration may become democratically destabilizing liabilities for the next.

A related and increasingly visible concern arises when public officials themselves are covertly influenced, improperly aligned, or quietly acting on behalf of foreign powers or concentrated private interests. In such cases, agentic systems wired into the machinery of government may gradually inherit and operationalize those hidden alignments at administrative scale, amplifying invisible human influence relationships while simultaneously making democratic responsibility harder for citizens to trace, understand, and publicly supervise. In addition, opaque proprietary systems, institutional

dependency, and vendor-controlled optimization architectures can further limit meaningful public visibility, making it harder to distinguish ordinary commercial optimization from the gradual emergence of civic influence.

The people building and promoting these systems are not necessarily malicious. Many are attempting, in good faith, to solve real societal problems.

But good intentions do not preserve constitutional limits.

Institutional momentum does not guarantee democratic restraint.

And systems optimized primarily for speed, scale, prediction, and administrative efficiency do not naturally preserve the civic conditions upon which republican government ultimately depends: human judgment, visible accountability, public deliberation, constitutional restraint, and the moral responsibility carried by free citizens governing themselves. In a constitutional republic, some democratic friction is not failure. It is part of the infrastructure of legitimate self-government.

At present, the United States appears increasingly committed to a rushed deployment of Agentic AI. In other sectors, this may be helpful - but not in U.S. government. Agentic AI is already rapidly becoming part of the operational infrastructure of governance, administration, security, and public decision-making. This book does not deny that reality. Nor does it claim that agentic AI should be prevented from operating elsewhere in society.

It does insist, however, that Agentic AI must never become Civic AI operating near power. Systems capable of autonomous strategic action, adaptive optimization, operational drift, and increasingly opaque forms of machine-directed governance must not become the governing intelligence of the American republic itself.

Democratic restoration therefore requires removing Agentic AI from most government-related functions at all levels - except for our military and law enforcement agencies.

Civic legitimacy requires a different standard: refusal-based Non-Agentic Civic AI designed to illuminate public reality while remaining structurally incapable of partisan judgment, ideological steering, autonomous recommendation, or inherited political alignment. Here, refusal-based does not mean incapable. It means constitutionally constrained: designed to refuse delegated civic judgment, autonomous political authority, and the optimization of citizens themselves.

This is not a plan for smarter artificial rulers. Not a plan for replacing human conscience with optimized computation. Not a plan for quietly transferring democratic judgment into systems too complex, adaptive, opaque, or institutionally entrenched for ordinary citizens to meaningfully challenge.

Instead, the pages that follow propose a different kind of civic infrastructure: Non-Agentic Civic AI designed not to govern human beings, but to help human beings govern themselves more clearly. This alternative framework is visibility-first, not automation-first. Its purpose is not to replace human judgment, but to preserve and strengthen it. Its aim is to illuminate civic reality while keeping responsibility, legitimacy, and final authority fully human. Where democratic government succeeds,

responsibility must remain human. Where democratic government fails, responsibility must remain human there as well.

A free people must never lose the ability to identify who governs, who decides, who bears responsibility, and who may be held accountable before the public.

You do not need to understand the technical mechanics of artificial intelligence to understand the central question this book asks: If Agentic AI in government is structurally prone to drift beyond meaningful human oversight, and if democratic responsibility cannot legitimately be delegated to artificial agents, are we wise to continue embedding such systems within the exercise of public authority without developing a more durable alternative?

This volume is my answer to that question.

It is an attempt to prepare, in advance, a workable democratic alternative before some future crisis forces the nation to confront these questions all at once.
Perhaps one morning, after a disturbing headline, institutional scandal, or visible failure of accountability, Americans will begin asking: How did we become so dependent upon systems we do not truly govern? Was there another path? What would democratic restoration require now?
My hope is that this work may help answer those questions.

Throughout this work, you will encounter terms that may initially seem unfamiliar: “Non-Agentic Civic AI,” “Civic Digital Visibility First,” “The 52-Gates Democratic Permissions Threshold,” and others.

These terms were not created merely as academic labels or proprietary concepts. They emerged gradually from the constitutional problems this work attempts to describe - problems for which our existing civic and technological vocabulary may not yet provide sufficiently precise public language. Sometimes new terms are needed when emerging technologies create distinctions older vocabularies cannot adequately express. Some readers may find these terms useful for public dialogue; others may prefer different language entirely. What matters is not whether these exact terms endure, but whether we find effective ways to approach the critical questions they are created to address.

For example, we must debate and decide: What are the constitutional boundaries between systems that assist human judgment and systems that increasingly substitute for it? Any successful transition to our Age of Humanity with AI will require an answer to this question. Free societies must preserve a meaningful boundary between systems that help citizens see and systems that begin quietly shaping, filtering, or replacing public judgment itself.

At the time of this writing, a growing number of institutions and researchers in fields such as constitutional law, AI safety, public administration, and cybersecurity have begun openly to confront a similar question: Where should the boundary lie between machine-assisted analysis and delegated machine judgment in public life?

Policy institutes, democracy advocates, technical governance bodies, public institutions, and AI developers have each, in their own language, begun describing the risks of allowing autonomous systems to assume roles once reserved for accountable human decision-makers.

No durable public consensus has yet emerged about how far such delegation should be permitted to go, or what constitutional limits a free society ought to place upon Agentic AI in government. *But we can all agree that* ***democracies cannot govern what they cannot see.*** Any serious democratic debate regarding boundaries must also confront a prior requirement: boundaries that cannot be clearly seen cannot be meaningfully compared, publicly judged, or constitutionally restrained. To deliberate responsibly about agentic boundaries, citizens, governments, institutions, and AI designers alike must first insist upon sufficiently shared visibility within our civic digital field. Before democratic alignment regarding such systems becomes possible, the boundaries themselves must become visible enough for a free people to understand, debate, and govern together.

The distinction between civic visibility infrastructure and delegated machine governance is not merely semantic. It is increasingly a structural question of democratic legitimacy.

Democratic society has always rested upon a simple understanding: no individual exists above the constitutional conditions that preserve the common good. Citizens accept limits not because liberty is weakened by restraint, but because freedom survives only when power remains accountable to the whole. This work proposes extending that same principle to Civic AI. If new forms of intelligence are to serve democratic society legitimately, they too must exist within constitutional boundaries designed not for domination, but for service. None of us are greater than all of us.

Agentic AI in U.S. Government: Optimization or Stranded Democratic Capital?

One of the most practical democratic questions raised by the rushed deployment of Agentic AI in American government is not only whether such systems are innovative, but whether they are even durable. Even if many current officials, vendors, and reformers are acting with good intentions, public institutions may be investing billions of taxpayer dollars in architectures that prove legally unstable, constitutionally contestable, and institutionally difficult to sustain over time.

The modern case for Agentic AI in government is usually framed in the language of optimization, efficiency, modernization, and administrative speed. Yet democracies must ask a prior question before scaling such systems: Can a republic safely keep what it is being asked to buy? If a public technology later proves incompatible with due process, accountable human judgment, transparent administration, or constitutional limits on delegation, then the original expenditure may no longer appear as prudent modernization, but as a costly and avoidable misallocation of public funds.

This is where the danger of what may be called ‘stranded democratic capital’ becomes clear. The problem is not that government spent money on technology. The problem is that it may spend large public sums on systems that democracy itself cannot safely or legitimately keep - systems that later require suspension, litigation, redesign, or removal once their civic incompatibilities become more widely understood. Public officials possess a fiduciary duty not merely to pursue innovation, but to avoid binding taxpayers to governance architectures that a constitutional republic may ultimately prove unable to sustain. In that event, taxpayers may ultimately pay twice: first to install fragile Agentic AI infrastructures, and later to dismantle or replace them.

Recent public-sector studies already suggest reason for caution. Governments are rapidly expanding advanced Agentic AI use even while lagging behind on trustworthy governance foundations, transparent oversight structures, and durable accountability mechanisms. Other analysts have warned that agentic systems can encourage administrative rubber-stamping, diffuse responsibility across

opaque algorithmic processes, and insert poorly understood forms of delegated power into institutions that remain constitutionally answerable to the public. If such concerns continue to deepen, many current investments may eventually come to resemble not durable civic infrastructure, but temporary experiments built upon democratically unstable ground.

The deeper issue is not technological capability alone, but whether democratic institutions can sustainably survive the forms of delegation such systems encourage. Throughout history, republics have repeatedly discovered that not every system that increases administrative efficiency is compatible with long-term civic legitimacy. Public institutions must therefore distinguish between technologies that strengthen democratic visibility and those that gradually displace accountable human judgment behind increasingly opaque systems of optimization.

A wiser course remains available. Democracies should invest in Non-Agentic Civic AI - civic visibility tools designed to illuminate conditions, strengthen shared public legibility, and support human judgment without replacing it. Such systems are far more likely to remain inspectable, lawful, publicly intelligible, and fiscally durable across generations. If democratic institutions are to modernize, they should do so in ways they can still justify, govern, audit, and afford decades later. In the end, the final measure is not whether governments can build increasingly agentic systems. It is whether free societies can sustainably absorb them without undermining the constitutional trust, accountability, and human responsibility upon which democratic legitimacy ultimately depends.

As you consider the ideas herein, please remember:

Our democratic legacy allows us to question whether every capability that can be deployed in government should be deployed.

We are allowed to insist that constitutional self-government requires more than optimized administration.

We are allowed to demand systems that illuminate our civic field rather than quietly govern it.

We are allowed to require that where government acts, responsibility remains human and visible.

We are allowed to ask whether using taxpayer funds toward developing and installing Agentic AI under increasing democratic, legal, and political scrutiny is wise, given our current non-agentic alternatives.

And we are allowed, together, to restore democratic responsibility before it disappears beneath the logic of efficiency entirely.

The pages that follow are an invitation to do exactly that.

Our future need not be quietly decided for us by others.

A Principled Call to Ban All Agentic AI from U.S. Government (All Levels)

I. The Temptation of Efficiency

If machines can drive cars, trade stocks, and write code, why shouldn't they help run government? For many Americans, that question does not feel theoretical. It feels like common sense born of frustration. They wait on hold with public offices. They navigate forms that seem designed to confuse. They watch crises unfold faster than agencies can respond and wonder why a country that can route millions of packages overnight struggles to process a permit in less than six weeks.

Under those conditions, the promise of agentic artificial intelligence in governance can sound almost merciful. What if digital "coworkers" could clear backlogs, route cases, and keep offices open 24 hours a day? What if proactive systems could fill out forms automatically, check eligibility in real time, and nudge people toward benefits they might otherwise miss? What if optimizers could scan thousands of projects and tell us, without politics, which roads, schools, or bridges should be fixed first?

It is important to take that hope seriously. Democracies do not win trust by romanticizing delay or defending every inefficiency as a virtue. In hospitals, courts, and welfare offices, slow decisions can mean real harm. In an age of tight budgets and rising expectations, public servants themselves often long for tools that could relieve them of repetitive work and let them focus on what truly requires human judgment.

Yet there is a reason earlier generations did not define a republic as "a government that does things quickly." Democracies are not designed to maximize efficiency. They are designed to limit and balance power. Procedures, debate, and even a certain amount of friction are not accidents in a constitutional order. They are how a free people keeps authority visible, contestable, and answerable for its use. Speed can be welcome. But when speed erases visibility, the cost is legitimacy. A free people does not merely ask how quickly power moves, but how clearly it can be seen.

What begins as assistance rarely remains so. Systems that help us decide have a quiet tendency, over time, to begin deciding for us. In the chapters that follow, this book will argue for a new class of Civic AI tools. They are non-agentic instruments of sight, not engines of decision. Before that argument can stand, however, we must confront an uncomfortable possibility. If we invite agentic systems into the machinery of government in the name of efficiency, they will not simply make old processes faster. Over time, they will change what it means to decide at all. That is the danger this chapter examines - not as a matter of misuse, but as a matter of structure.

II. Agentic Systems and the Civic Ceiling

Before we can assess the dangers of agentic AI in governance, we must be precise about what we mean by "agentic" and about the line this book treats as constitutionally non-negotiable. Vagueness is itself a risk here. If every automated spreadsheet is called "agentic" or every database query is called "Civic AI," then the terms lose the very clarity democracies need when new powers approach public authority.

In this work, a system counts as agentic when it does more than represent, retrieve, or illuminate information. It is agentic when it can, on its own initiative and within a defined scope, select actions, change states, or steer processes without fresh human judgment at each consequential step. An agentic system may open a case, close an application, route a file, trigger an alert, or assign a rank that other parts of government then treat as operative. Its defining feature is not intelligence, but delegated discretion. It is allowed to move.

By contrast, this book reserves the term Civic AI for a far narrower and more disciplined class of tools. Civic AI, as defined here, is structurally non-agentic. It may extend human perception, assemble and cross-reference records, detect patterns, and, when explicitly asked, provide analysis that clarifies options or consequences. It may clarify the consequences of choices, but it may never substitute for the act of choosing. Its vocation is illumination, not rule.

This distinction implies a civic ceiling - a boundary beyond which no machine may exercise practical discretion without eroding the conditions of self-government. Below that ceiling lie domains of function: logistics, formatting, internal scheduling, technical monitoring, and other forms of mechanical assistance that do not, in themselves, determine who is heard, who is served, or who is constrained. Above that ceiling lie domains of judgment: rights, eligibility, enforcement, durable access to public goods, and the structuring of speech, participation, and political choice.

The central claim of this chapter is simple. Systems that cross this civic ceiling - by deciding who qualifies, who is prioritized, who is watched, or how coercive power is applied - may be technologically impressive and administratively useful. They may be discussed as tools of public administration. But they cannot rightly be called Civic AI, nor can they be admitted near democratic authority without displacing it. Civic AI, by definition in this framework, occupies only one of three roles in a constitutional order: humans judge, bounded systems execute narrow functions below the civic ceiling, and Civic AI helps all parties see what is happening in time to decide. Anything more is not illumination; it is governance.

If this line feels severe, the rest of the chapter exists to test it. We will not ask whether agentic approaches can be built, or whether they might increase efficiency in some measure. We will ask whether, once admitted across the civic ceiling, they can avoid altering who effectively governs. Under that test, we will find that the danger is not occasional misuse, but structural drift.

III. The Structural Drift Principle

The risks of agentic systems in governance do not arise only from bad actors, biased data, or flawed implementation. They arise from a quieter and more stubborn pattern in how delegated systems evolve over time. This pattern can be stated simply:

- Systems that assist begin to recommend.
- Systems that recommend begin to be followed.
- Systems that are followed become embedded.
- Systems that become embedded begin to govern.

At the outset, an agentic tool may be introduced only to "help" human decision-makers manage complexity or volume. Its outputs are framed as suggestions, its role as advisory. But as time passes, two forces converge. First, institutions under pressure learn to rely on whatever reduces workload and controversy. Second, the very existence of a system that appears more consistent, more data-driven, or more objective makes it harder to depart from its guidance without justification. Deference becomes the norm, even when it is not formally required.

Once deference hardens into habit, the system's role changes. What began as one input among many becomes the practical baseline. Human judgment still exists, but it operates within a narrowed corridor defined by the machine's categories, rankings, or alerts. Over time, organizations are re-structured around these defaults: workflows, oversight routines, and even legal defenses begin to presume that "following the system" is the ordinary, prudent course. At that point, the tool is no longer merely assisting authority. It is helping to constitute it.

This book calls that progression structural drift. It does not depend on ill will. It is the expected behavior of delegated systems operating under institutional pressure. In private life, such drift may be unfortunate. Near democratic authority, it is disqualifying. Any agentic system allowed to operate across the civic ceiling will, if it proves even modestly useful, tend over time to move from support to substitution. That is why, in this framework, non-agentic design is not a preference. It is the constitutional precondition for anything that wishes to be called Civic AI.

IV. The Emerging Ecosystem of Agentic Governance Systems

The Structural Drift Principle would matter little if agentic systems near public authority were only speculative. They are not. Across the United States, a recognizable ecosystem of approaches is already being proposed, piloted, and in some cases deployed in and around government. They are described in hopeful language - GovAI, digital government, AI-enabled public services, agentic workflows - but structurally they share one feature: delegated discretion operating close to the civic ceiling.

For clarity, this chapter will group these efforts into six families. The labels are analytic, not official. They are offered to show that behind varied branding and technical detail, a small number of recurring patterns are emerging.

1. Automated Eligibility and Risk Systems:

The first family consists of systems that determine or heavily influence who qualifies, who is flagged, and who receives. These tools use past data to generate scores or classifications that guide decisions about individuals: eligibility for benefits, perceived risk in criminal justice or child welfare, likelihood of fraud in public programs, or priority in inspections and enforcement.

In formal terms, such systems transform data into discrete categories - approved or denied, low risk or high risk, routine or exceptional - that other institutional processes then treat as operative. Even when a human remains "in the loop," the default path is clear: the score is presumed correct unless overridden. Delegated discretion resides not only in explicit auto-approval or auto-denial, but in the way scores quietly reset the starting point of judgment. The system has redrawn, in practice, who the state sees as trustworthy, deserving, or dangerous.

2. Resource Prioritization and Planning Optimizers

A second family focuses less on individuals and more on where public resources go first. These are optimization and planning systems that rank projects, communities, or interventions according to specified objectives. They may seek to maximize cost-effectiveness, reduce measured disparities, improve resilience, or satisfy composite indices of "need" or "impact."

Here, the agentic element lies in the production of ordered lists or recommended portfolios that quickly become the default structure for budgets and plans. Once accepted as the rational way to allocate scarce resources, these rankings function as de facto policy. The choice of objective function - what to maximize, what to minimize, what to count and what to ignore - is itself a political judgment. Yet after the function is fixed, its outputs present as technical necessity rather than as the contested trade-offs they truly encode.

3. Agentic Workflows and Digital Coworkers

A third family operates inside agencies under the banner of workflow automation and "digital coworkers." These systems do not only store or summarize information. They route cases, open and close tasks, trigger follow-ups, escalate files, and in some instances move multi-step administrative processes forward with minimal human intervention. Their purpose is to clear backlogs, standardize handling, and ensure that no case "falls through the cracks."

In this domain, delegated discretion appears in timing and sequence. Who is contacted first, which applications are flagged as incomplete, which files are quietly closed after a threshold of non-response - all become emergent properties of routing logic and system rules. Over time, the practical experience of a right or benefit - how many hurdles must be cleared, how long a person must wait, how likely a case is to be revived - depends less on formal law than on how these internal agents manage queues and thresholds.

4. Autonomous Citizen-Service Agents

A fourth family sits directly at the interface between citizen and state: autonomous or semi-autonomous service agents. These systems go beyond answering questions or providing information. They fill out forms, check eligibility against existing records, submit applications, open cases, schedule appointments, and sometimes resolve service requests end-to-end. Some are explicitly designed to act "on behalf of" the resident.

Here, discretion is exercised in the interpretation of intent and in the construction of a path. The agent decides what the person meant to do, which programs or remedies are most relevant, which steps are necessary, and when the interaction is complete. Participation shifts from a citizen taking public actions in their own name to a system acting as their procedural surrogate. Misunderstandings, omissions, and biases in this surrogate's behavior may be hard to detect, especially for those with the least capacity to question the outcome.

5. Autonomous Security, Fraud, and Cyber Agents

A fifth family is justified in the language of protection: cybersecurity agents, fraud-detection engines, and anomaly-response systems. These tools monitor networks, transactions, and flows of activity continuously. They are designed not merely to log anomalies, but to act on them - to isolate machines, block accounts, trigger holds on payments, or open targeted reviews, often at machine speed.

In these settings, discretion is expressed through thresholds and automated responses. The system decides what counts as suspicious, when an alert warrants containment rather than observation, and how much collateral disruption is acceptable in the name of safety. Human review may occur after the fact, but the initial move - lock, freeze, flag, or sever - has already changed the practical situation of those affected. "Better safe than sorry" ceases to be an explicit human posture and becomes the standing behavior of a machine.

6. Advisory-to-Agentic Governance Tools

The final family begins life in the least threatening form: advisory and analytic systems that help policymakers or administrators make sense of complex evidence. These tools synthesize documents, model scenarios, compare options, and generate ranked lists of policies or actions. At first, they are framed as decision-support - another lens through which humans can view the field.

Over time, however, advisory systems tend to move along a familiar spectrum, following the same pattern described earlier: from assistance to recommendation, from recommendation to default, and from default to embedded governance. Recommendations become defaults. Defaults become automated in routine cases. Exceptions become rare. The tool that once merely described or suggested quietly shapes which options are even considered and what counts as a "reasonable" choice. In this way, design decisions about what is modeled, which values are encoded, and how options are ranked begin to steer the civic agenda itself.

7. Review of the Emerging Ecosystem of Agentic Governance Systems

Taken together, these six families constitute the practical field this chapter addresses. They differ in domain and rhetoric, but they share two structural traits: they operate close to or across the civic ceiling, and they embody the pattern of structural drift. They differ in function, but converge in consequence: the quiet relocation of judgment from human institutions to system behavior. Each promises assistance. Each is built to be useful. And, as the next section will argue, each contains within itself a characteristic way of moving from support to substitution that no democracy can safely ignore.

Taken on their own, these six families might still seem like a loose collection of innovations - each with its own promise, each with its own risks to be managed. What the Structural Drift Principle reveals is that they share a common trajectory. Each begins as assistance at or above the civic ceiling. Each embeds delegated discretion into decisions about who counts, who is served, who is watched, or which options are even considered. And each, if it proves useful, is pulled by ordinary institutional pressures toward becoming the place where practical judgment actually occurs.

Seen together, these six failure modes are not a scatter of unrelated problems. They are six variations on one theme: authority sliding, step by step, from visible human institutions into system behavior that is hard to see, harder to contest, and hardest of all to reverse once organizations have grown around it. That is structural drift in civic form. And because each family crosses or presses against the civic ceiling - exercising delegated discretion in domains of rights, eligibility, enforcement, access, and agenda-setting - each is, within this framework, constitutionally inadmissible as Civic AI, however effective or well-intentioned it may appear.

AGENTIC AI—FLOW OF FAILURE in AMERICAN GOVERNANCE

Across the United States, six families of agentic systems are emerging in and around governance. Each promises efficiency; each contains a characteristic democratic failure.

First, **automated eligibility and risk systems** promise faster, more consistent, and "objective" decisions about who qualifies, who is flagged, and who receives public benefits or scrutiny. They use past data to generate scores that other processes then treat as operative. The structural failure here is a hidden re-weighting of who counts: scores quietly redraw the civic boundary of trust and standing, shifting who is treated as deserving or dangerous from visible law and debate into opaque model behavior. In doing so, they breach equal civic standing and invite the moral scoring of persons.

Second, **resource prioritization and planning optimizers** promise "rational," data-driven allocation of limited funds and projects to maximize impact, efficiency, or fairness across communities. They generate ordered lists and recommended portfolios that quickly become the default structure for budgets and plans. Their structural failure is optimization without representation: political trade-offs are encoded in objective functions and then presented as technical necessity, so de facto policy is made inside algorithms rather than in legislatures. Human judgment over ends is displaced, and contested value choices are quietly outsourced.

Third, **agentic workflows and digital coworkers** promise reduced backlogs, standardized handling, and seamless case management through automated routing, escalation, and closure. They route cases, open and close tasks, and advance multi-step administrative processes with minimal human intervention. Their structural failure is procedural drift into invisible policy: the lived experience of rights and benefits - delays, hurdles, odds of revival - is quietly rewritten by workflow logic never authorized as law. Decisions are effectively made inside process code, eroding visibility of mandate and accountability.

Fourth, **autonomous citizen-service agents** promise frictionless, end-to-end access to services by allowing systems to act "on behalf of" residents - completing forms, checking eligibility, and submitting requests. Here the structural failure is substitution of voice: participation is mediated by a procedural surrogate that interprets intent and chooses paths, so direct civic agency is replaced by system behavior that is hard to see or contest, especially for the least powerful. Human agency in civic acts is displaced, and contestability is weakened.

Fifth, **autonomous security, fraud, and cyber agents** promise continuous protection of infrastructure and funds through real-time anomaly detection and automatic containment, holds, or targeted reviews. They monitor networks and flows of activity and act on anomalies at machine speed. Their structural failure is emergency logic without politics: locking, freezing, and flagging become standing machine behaviors rather than explicitly owned human choices, normalizing quasi-coercive actions without visible warrants of judgment and straining rights protections under stress.

Sixth, **advisory-to-agentic governance tools** promise better policymaking through summaries, simulations, and ranked options that make complex evidence more tractable and "evidence-based." They begin as decision-support systems but tend to move from assistance to recommendation, from recommendation to default, and from default to automation. Their structural failure is soft capture of the agenda: recommendations slide into defaults and defaults into embedded governance, so the real power lies in what is modeled and ranked, while elected officials govern within a pre-framed corridor of "reasonable" options. Agenda-setting is quietly outsourced, and pluralism narrows.

The following map summarizes that pattern. For each family, it sets out the efficiency claim that makes the system attractive, the named failure mode that emerges as drift takes hold, and the underlying civic principle strained or breached. All six operate where delegated discretion and democratic judgment meet. The point is not that every deployment will immediately collapse these safeguards, but that the structure of each approach makes the same kind of failure more likely the longer it remains in place.

AGENTIC AI - FLOW OF FAILURE in AMERICAN GOVERNANCE (Chart)

All six families operate at or across the civic ceiling, where delegated discretion begins to displace visible human judgment.

Agentic system family	Efficiency claim (the temptation)	Structural failure (democratic break)
Automated eligibility & risk systems	Faster, more consistent, and "objective" decisions about who qualifies, who is flagged, and who receives public benefits or scrutiny.	**Hidden re-weighting of who counts.** Scores quietly redraw the civic boundary of trust and standing, shifting who is treated as deserving or dangerous from visible law and debate into opaque model behavior. This breaches equal civic standing and invites moral scoring of persons (Gates 4, 27, 33).
Resource prioritization & planning optimizers	"Rational," data-driven allocation of limited funds and projects to maximize impact, efficiency, or fairness across communities.	**Optimization without representation.** Political trade-offs are encoded in objective functions and then presented as technical necessity, so de facto policy is made inside algorithms rather than in legislatures. This displaces human judgment over ends and outsources contested value choices (Gates 11, 14, 26).
Agentic workflows & digital coworkers	Reduced backlogs, standardized handling, and seamless case management through automated routing, escalation, and closure.	**Procedural drift into invisible policy.** The lived experience of rights and benefits—delays, hurdles, odds of revival—is quietly rewritten by workflow logic never authorized as law. This obscures where decisions are really made and erodes visibility of mandate and accountability (Gates 3, 15).
Autonomous citizen-service agents	Frictionless, end-to-end access to services by letting systems act "on behalf of" residents to complete forms, check eligibility, and submit requests.	**Substitution of voice.** Participation is mediated by a procedural surrogate that interprets intent and chooses paths, so direct civic agency is replaced by system behavior that is hard to see or contest, especially for the least powerful. This displaces human agency in civic acts and weakens contestability (Gates 5, 25, 42).
Autonomous security, fraud, and cyber agents	Continuous protection of infrastructure and funds through real-time anomaly detection and automatic containment, holds, or targeted reviews.	**Emergency logic without politics.** Locking, freezing, and flagging become standing machine behaviors rather than explicitly owned human choices, normalizing quasi-coercive actions without visible warrants of judgment. This violates the refusal of autonomous coercive action and strains rights protections under stress (Gates 24, 30).
Advisory-to-agentic governance tools	Better policymaking through summaries, simulations, and ranked options that make complex evidence tractable and "evidence-based."	**Soft capture of the agenda.** Recommendations slide into defaults, and defaults into automation; over time, the real power lies in what is modeled and ranked, while elected officials govern within a pre-framed corridor of "reasonable" options. This quietly outsources agenda-setting and narrows pluralism (Gates 26, 48).

V. Efficiency and Legitimacy

To this point, the analysis may seem to place efficiency and democracy in simple opposition, as if anything that speeds public work must be suspect. That is not the claim. In many settings, better tools, clearer information, and thoughtful automation can ease genuine burdens. A benefits office that reduces paperwork, a court that shortens needless delays, or a city that coordinates repairs more smoothly may all serve justice rather than weaken it. The six families described here are compelling in part because they address real institutional strain.

The difficulty is not that these systems fail to deliver gains. It is that they deliver those gains by moving discretion upward, across the civic ceiling, into opaque machinery that is no longer answerably human. Structural drift converts efficiency gains into a quiet re-location of judgment from visible institutions to embedded defaults in code and workflow. The more an agency depends on such systems for basic functioning, the harder it becomes in practice to refuse their outputs, contest their assumptions, or dismantle them without crippling essential services. What began as a choice to adopt a tool hardens into a dependency that effectively governs. At that point, efficiency is no longer simply a feature of governance; it has become a pathway to its quiet redefinition.

Democracies cannot accept that trade. Their first obligation is not to do everything faster, but to keep authority anchored where it can be seen, challenged, and owned. Legitimacy rests on the assurance that consequential decisions about rights, access, and coercive power are made by people who can be named, questioned, and, if necessary, replaced - not by systems whose inner workings few can explain and no one person is responsible for. When agentic designs cross the civic ceiling, they offer a bargain: fewer backlogs, smoother surfaces, quieter conflict, in exchange for a slow blurring of who actually decides. Under this framework, that bargain must be declined. Non-agentic Civic AI may help human institutions see more clearly and move more promptly. It may not, and must not, be replaced by systems that decide in their place.

On “Limited Agency,” “Bounded Autonomy,” and “Human-in-the-Loop”

Some will reply that “limited agency” is enough - that if systems act only within narrow bands, the danger is reduced to a manageable level. But structural drift does not depend on breadth of mandate. Even narrowly scoped systems, once embedded around recurring judgments about who qualifies, who is prioritized, or who is watched, begin to define the practical meaning of those judgments. The question is not how large the zone of agency is, but whether any zone of delegated judgment has been admitted across the civic ceiling at all. Once that line is crossed, the same progression from assistance to governance still applies.

Others will argue for “bounded autonomy” - agentic systems carefully constrained by rules, safeguards, and oversight. Yet from a constitutional perspective, this simply relocates the problem. If a system can open cases, close applications, rank persons, or trigger coercive steps without fresh human judgment in each instance, it is exercising authority in domains that a republic has no warrant to delegate to machines. The presence of guardrails may reduce harm; it does not change the fact that decisions once made under visible law and accountable officials are now being made inside code and control logic.

A third reassurance is that "human-in-the-loop is enough." In practice, however, structural drift turns many such humans into ratifiers rather than deciders. When institutional routines, legal defenses, and performance metrics all treat "following the system" as the prudent default, the burden of proof shifts onto anyone who would depart from its outputs. Over time, the loop reverses: it is the machine that effectively sets the baseline, and the human who must justify exceptions. Formal presence in the loop does not guarantee that judgment still resides with people; it often masks the fact that authority has already migrated into the system's defaults.

For that reason, the issue is not how agency is limited, bounded, or overseen, but whether it is admitted at all in domains where judgment must remain visibly human. In such domains, the question is one of admissibility, not calibration.

VI. Constitutional Verdict: Agentic AI as Power Near Democracy

By this point, the pattern is clear. Agentic systems that operate at or above the civic ceiling do not merely assist democratic governance. Over time, they reconfigure it. Automated eligibility and risk engines, optimization tools, workflow agents, citizen-service surrogates, autonomous security systems, and advisory-turned-agentic platforms each embody the same structural drift: from assistance to recommendation, from recommendation to default, from default to embedded governance. However different their technical details, they all move practical judgment away from visible human institutions and into system behavior that is difficult to see, contest, or reverse.

Within the framework of this book, that pattern is not a technical flaw to be tuned away. It is a constitutional failure. Civic AI, as defined here, begins with constraint rather than capability. It is non-agentic by design. It may illuminate conditions, assemble records, and analyze patterns on request. It may never decide who qualifies, who is prioritized, who is watched, how coercive power is applied, or which civic options count as "reasonable." Those are domains of judgment, and judgment in a republic must remain visibly and answerably human. Any system that crosses this line, however efficient or well-intentioned, stands outside the Civic AI category. It is simply powerful administrative AI operating near democracy.

Applied to the six families mapped in this chapter, the constitutional verdict follows directly. Automated eligibility and risk systems fail by redrawing civic standing through opaque scoring. Resource optimizers fail by encoding political trade-offs as technical functions. Agentic workflows fail by rewriting the lived content of rights through unvoted procedures. Citizen-service agents fail by substituting a procedural surrogate for the citizen's own civic voice. Autonomous security and fraud agents fail by normalizing machine-driven quasi-coercive action. Advisory-to-agentic tools fail by capturing the agenda within a pre-framed corridor of options. Each breaches, in its own way, the separation of illumination and authority that the civic ceiling is meant to enforce.

This does not mean that a democracy can never use automation or advanced computation in administration. It means that when such systems are agentic in domains of judgment, they may not be treated as Civic AI, and they may not be admitted near public authority under the standards proposed here. They may be regulated as ordinary administrative tools, scrutinized as sources of risk, or rejected as incompatible with self-government. But they cannot share in the name or vocation of Civic AI, whose sole legitimate role is to extend sight without sharing in rule.

The work of the next chapter begins there. If agentic AI must be treated as power near democracy rather than as Civic AI within it, then a different kind of system is needed to watch the edge where that power operates. Non-agentic Civic AI, properly bounded, can serve as a lantern at that boundary - monitoring structural drift, illuminating where authority is quietly shifting, and buying back time for human institutions to deliberate. It cannot co-govern. It can help a republic keep its own promises in the presence of new and tempting tools.

An American republic, Benjamin Franklin once warned, would endure only "if you can keep it." The agentic systems described in this chapter do not announce themselves as threats to that charge; they arrive as helpers, optimizers, and digital coworkers. Yet their structure pulls authority away from the visible, answerable human judgment that keeping a republic requires. Civic AI, in this framework, is permitted only one role: to help a free people see what they are doing in time to decide for themselves. Anything more, however sophisticated, belongs on the other side of Franklin's warning - not as an instrument of keeping the republic, but as one more way it can quietly be lost.

Watching the Edge: Civic AI and the Monitoring of Agentic AI Drift

As artificial intelligence becomes more capable, it is unsurprising that governments will be tempted to use increasingly autonomous, "agentic" systems to manage public work. Routing vehicles, scheduling repairs, adjusting energy loads, forecasting demand - these are the kinds of functional tasks that lend themselves to automation. When such systems remain within the domain of logistics and engineering, they raise familiar questions of safety, reliability, and procurement. They do not, in themselves, constitute Civic AI.

The danger arises when tools built for function begin, often quietly, to exercise something closer to authority.

A routing system that merely optimizes bus schedules is one thing. A routing system that gradually withdraws service from certain neighborhoods - converting efficiency into exclusion without visible public decision - is another. In practice, many of the most consequential shifts in democratic life do not arrive as laws or doctrines. They arrive as "upgrades," "optimizations," and "recommendation engines" that reshape who can participate, who is heard, and who is left waiting at the curb.
If agentic systems are permitted anywhere near the machinery of government, even at its logistical edges, a republic must cultivate a new form of vigilance: the capacity to see when such systems begin to move inward. This is one of the central civic responsibilities of non-agentic Civic AI.

Civic AI does not co-govern. It does not tune, direct, or control. Its vocation is simpler and more constitutional: to watch the edge of authority.

Within this framework, Civic AI carries three linked responsibilities.

Monitoring Duty: Civic AI observes the aggregate behavior and outputs of agentic systems operating within or alongside public institutions. This observation is conducted strictly at the level of systems and patterns, never at the level of individuals. It tracks how such systems redistribute access, burden, and institutional dependence over time - for example, whether automated scheduling systems consistently reduce service to particular districts, or whether triage tools shift waiting-time burdens onto specific populations.

Boundary-Watch Duty: Civic AI detects when these patterns approach the line of civic judgment. That line is crossed whenever a system begins, in practice, to influence rights, eligibility, enforcement, speech conditions, or durable access to core public goods such as education, justice, health, or democratic participation. What appears as neutral optimization can, over time, become a de facto allocation of civic standing.

Alert-Only Duty: When such drift is detected, Civic AI performs one function only: illumination.
It issues clear, non-binding notices to named human authorities and oversight bodies, explaining what has changed, why the shift may constitute movement toward civic authority, and which questions now require human deliberation. It does not recommend remedies, rank outcomes, or trigger automated responses. Machines may reveal that a boundary is being tested. Only people may decide how to respond.

This role must remain structurally bounded.

Civic AI may observe agentic systems, but it may not interact with, direct, or modify them. It must remain outside any operational control loop. The moment it begins to act upon the systems it observes, it ceases to be a Civic AI and becomes another agent within the system it was designed to illuminate. This constraint is not a limitation. It is the source of its legitimacy.

Agentic systems share a structural tendency: they do not remain static. Systems designed to assist begin to recommend. Systems that recommend begin to be followed. Systems that are followed become embedded. And systems that become embedded begin, often invisibly, to shape outcomes that were once the product of human judgment.

This progression is rarely declared. It emerges gradually - through convenience, efficiency, and institutional adaptation. Its effects, however, are cumulative. What begins as support can become substitution.

A republic cannot rely on intention alone to guard against this transformation. It must be able to see it. Civic AI provides that sight.

If agentic systems are to exist at the margins of governance, then a republic worthy of the name must also build instruments capable of seeing when those systems begin to move inward. Without such visibility, boundaries erode silently, authority shifts without acknowledgment, and democratic accountability weakens without clear cause.

With visibility, the boundary remains real.

The principle is simple: Delegation may be acceptable in domains of function. It is not acceptable in domains of judgment.

Within democratic governance, every consequential decision about who counts, who is heard, and who belongs must remain visibly and answerably human.

Civic AI helps ensure that this line remains clear - not by holding power itself, but by making power's movements visible in time to be governed.

It watches the edge.

Chapter 6: Our Non-Agentic Alternative:

- **"Civic Digital Visibility First – The Governing Principle of Non-Agentic Civic AI (For All U.S. Civic AI at All Government Levels"**

Civic Digital Visibility First: The Governing Principle for All Truly Civic AI

(Ohio Academy of History Speech - Marietta College - March 21, 2026)

(First Published at Medium.com on Mar 8, 2026)

The federal government and many state and local government offices are now actively encouraging the use of AI in government. That may be well-intentioned but it is risky.

Currently, about 85 percent of government offices rely on a single word processing software. Let's imagine a hypothetical scenario where a company chose to introduce a new product: "*Civic Word: Now with AI that can prioritize government projects!*" As a manager, I think, "I can call my staff together for a long, heated debate about our projects, or I can just push a button and let Civic Word rank them for us." At that moment, a democratic decision that once required human debate quietly becomes a software feature.

American democratic government - at the federal, state, and local levels - should reject all AI from government settings by default until that AI has been independently certified as compliant with democratic principles. We must not let 'optimizing' AI slip into our government sector on the strength of efficiency claims alone - Trojan Horses quietly treated as if they were Civic AI simply because they are deployed in civic settings.

We cannot jeopardize 2500 years of democratic evolution.

OUR RAPIDLY GROWING PROBLEM

Many assume that democracy is permanent - as if once established, it could only move forward, never backward. But we know that the historical record is not a straight line of democratic progress. It is a record of gains and losses - of republics that rose, stabilized for a while, and then often weakened or fell.

In recent years, we have watched signs of that weakening spread here in the United States and across the globe: eroding trust, captured institutions, elections held - but not believed.

These pressures did not begin with artificial intelligence. But AI - by neglect or by intent - can now accelerate these pressures at machine speed. It can deepen fractures faster than our civic efforts can repair them.

At the heart of this danger lies something very simple: democracies cannot govern what they cannot see.

So today I want to suggest that our central civic crisis is not a crisis of moral darkness. It is the fragmentation of the moral light that we still possess. But we're losing our ability to see clearly enough, quickly enough, and together enough to govern responsibly.

So how can democracies see clearly enough to govern responsibly in the digital age?

We must first understand the magnitude of the challenge. We can begin by looking at our own headlines. They reveal a pattern: how our individual and collective consciences are strained under a constant flood of crisis and confusion - and how that strain is beginning to erode the historically evolved civic processes on which democracy depends.

The pattern moves outward in six stages. Together these form what I call the *Six Rings of Digital Invisibility and Civic Destabilization*. Each ring widens the gap between democratic responsibility and democratic sight.

Ring One - Moral Overload

Modern citizens now absorb more suffering, outrage, and breaking news events in a single week than many earlier generations encountered in a year: war, climate disruption, political fracture, public heath crises, economic precarity, algorithmically amplified conflict.

The result is not moral indifference; it is moral saturation. Our sense of *ought* is triggered again and again - without resolution and without rest. People do not become immoral under that load; they become overwhelmed. Deliberation gives way to exhaustion. Outrage and withdrawal alternate. And when conscience is constantly triggered but rarely resolved, it slowly loses its capacity to guide.

Ring Two - Moral and Civic Invisibility

AI now operates inside financial markets at speeds regulators describe as "too fast to stop." Financial deepfakes have already moved markets with false announcements. Real harm occurs - but it unfolds across platforms, jurisdictions, and milliseconds. By the time we see the consequences, the chain that produced them is no longer visible.

For most of human history, harm was nearby. Suffering had faces. Responsibility could be traced. Today, harm increasingly occurs at a distance - mediated by platforms and fragmented across systems. People may sincerely wish to do no harm and yet participate, unknowingly, in systems that cause it. Conscience does not disappear here - but it loses its bearings.

Ring Three - Institutional Blindness

Modern institutions increasingly rely on systems they cannot fully see into. Courts, hospitals, regulators, election systems, and public-benefits programs now depend on digital infrastructures whose inner workings remain opaque.

We read about AI-assisted cyberattacks generating constantly mutating malware. We see attacks on hospitals, airports, and water systems. Government agencies deploy algorithmic tools whose logic even their operators cannot fully explain. Authority remains but comprehension thins. Institutions exercise power with diminishing visibility into how outcomes are produced. And when systems act beyond meaningful human oversight, legitimacy begins to erode.

Ring Four - Amplification Without Judgment

Deepfake videos circulate before verification can catch up. AI systems generate thousands of message variations to discover which narratives provoke the strongest emotional response. A growing share of political content online is synthetic or heavily manipulated.

Our digital systems are extraordinarily good at scaling what travels fastest - outrage, fear, simplified stories. But they are not designed to distinguish wisdom from distortion. When synthetic media becomes indistinguishable from authentic recording, shared reality begins to fracture. Democratic deliberation cannot function if it cannot discern.

Ring Five - Governance Under Strain

Across the United States and the world, governments are scrambling to respond to AI's public impact. Hundreds of legislative proposals appear. Executives invoke competitiveness. Legislatures demand oversight. Courts attempt to referee technical disputes with only partial visibility.

Leaders must make consequential decisions without a stable shared picture of what is actually happening. Government shifts from stewardship to permanent crisis response. Authority reacts to pressure rather than acting from understanding.

Ring Six - Cascading System Consequences

A ransomware attack in one country shuts down hospitals in another. Synthetic narratives created by small groups are amplified by foreign actors, domestic factions, and automated systems. A single misperception can move from social media to financial markets, to politics, to street-level conflict.

The rings reinforce one another: overload - invisibility - institutional blindness - amplification without judgment - governance under strain. By the time the damage becomes visible, it is already widespread and politically difficult to repair. Consequences now arrive faster than responsibility can be traced.

The Diagnosis

The *Six Rings* are not a prophecy of democratic collapse. They are a diagnosis of what our own headlines already describe: a progressive breakdown in our shared ability to see, interpret, and coordinate responsibly in a world moving at digital speed.

The problem is not a sudden loss of intelligence or goodwill. It is something more structural. Our civic environment is now moving faster than human perception was built to manage. We face a growing mismatch between moral responsibility and perceptual capacity.

And if this crisis is a collapse of civic visibility, then the central question becomes this: What could help our democracy see again - together - without giving our moral authority to machines?

Perhaps part of the response lies in something unexpected. If AI-acceleration is amplifying our problems, perhaps the development of democratically-certified Civic AI could become part of our solution - as a new instrument of sight.

Franklin's Experiment

It may help to remember an early American experiment. During a storm, Benjamin Franklin flew a kite carrying a key into a thundercloud and revealed the invisible electrical forces moving through the sky. He made an unseen force visible to human understanding. That new visibility soon produced a practical invention: the lightning rod - a tool of protection.

Properly designed, Civic AI could serve a similar function for democracy today. Imagine Civic AI acting like Franklin's kite and key for the digital age - drawing down hidden, harmful digital currents that affect civic life and making them visible to human discernment, while never becoming an instrument of power itself.

Civic AI can make democracies faster at seeing. Every moment saved in identifying hazards, tracing patterns, or assembling information at machine speed becomes one more moment available for human deliberation. More time. A clearer record. Better judgment.
But when designing Civic AI, we must ensure that it does not damage our democracy more than the problems it is meant to address.

A PROPOSED SOLUTION: DEMOCRATICALLY-CERTIFIED CIVIC AI

Democracies are not designed to maximize efficiency. They are designed to limit the concentration of power. And in the digital age, that means limiting the concentration of agentic capacity - systems that can act, steer, and scale influence.

AI is now being rapidly integrated into federal government, as well as state and local governments. The problem I see is this: AI does not become Civic AI simply because it is deployed in civic contexts. I propose that AI be banned from government use until it is proven to be democratically-certified Civic AI.

But before defining Civic AI, we must consider several fundamental questions:

- Under what constitutional conditions may new power approach public authority at all?
- Will civic visibility remain shared public infrastructure?
- Will civic visibility become concentrated in systems and institutions that see more than citizens do?

We don't need Civic AI to tell us what is right. We need it to tell us what is real. And so, in the years ahead, I suspect one principle will matter most: democratic illumination.

Not capacity. Not agency. Democratic illumination.

Moral visibility is the structured illumination of civic conditions where human responsibility is at stake - without directing, ranking, or resolving that responsibility. It makes consequences visible: who is affected, how, and over what time frame. It reveals patterns and chains of responsibility that might otherwise disappear inside digital complexity, and it gives human conscience something honest to look at. It does not preach or assign virtue, rank citizens, or engineer agreement. Moral visibility does not tell us what to believe; it simply makes the stakes harder to ignore.

First sight, then deliberation. Democracies already understand this sequence. Courts, legislatures, and voters seek information and clarity before acting. And yet, in many current discussions of AI, that sequence is quietly reversed. We begin with persuasion - optimization - mobilization. And visibility somehow becomes an afterthought.

Democracy works best in the opposite order: First sight, then deliberation. Because democracies cannot govern what they cannot see.

A properly ordered civic process might look like this:

- First, visibility - illumination of facts, hazards, and patterns.
- Second, analysis on request - synthesis, comparison, explanation.
- Third, human deliberation - citizens, journalists, courts, and legislatures weighing competing values.
- Fourth, persuasion and activism - if chosen.
- And finally, law and enforcement - if authorized.

Visibility does not compete with civic action. It makes civic action possible.

The premise of my *Civic Digital Visibility First* principle is simple. Treat Civic AI as visibility infrastructure - as digital lighthouses that warn us of digital hazards. Under this approach, a Civic AI system may illuminate hazards and inequities, reveal patterns and correlations, and provide analysis when requested.

But it must always stop short of the final sentence: "Therefore, you *should* do X." That final *should* belongs to human beings who must bear the moral and civic consequences of their choices.

Moral visibility is not a side effect here. It is the entire purpose. Civic AI should be a witness, not a sovereign. To maintain that boundary, we need explicit design constraints.

My work proposes what I call the *52-Gates Democratic Permissions Threshold*. Each gate represents a democratic principle and acts as a binary (yes/no) constitutional test for whether a given AI may qualify as Civic AI and be admitted near democratic authority.

Remember the Civic Word hypothetical scenario I opened with? Our *52-Gates* would have rejected "Civic Word" because it violates the tests of Illumination, Non-Agency, and Moral-Outsourcing (Gates 7, 4, and 12) - ensuring that project ranking remains a human debate, not a software feature.

No independent decision-making. No civic scoring. No person-level ideological ranking. No hidden authority. And if those boundaries cannot be maintained, the system's proper democratic response is refusal. Because the mission of Civic AI is not power; it is sight.

Not capacity. Not agency. Democratic illumination.

AI near democratic authority must be designed very differently from the systems we're often encouraged to imagine. Many people now hope for a kind of machine savior - a system that will decide faster, optimize better, persuade more effectively.

But optimization is not wisdom. A system optimized for speed can rush toward a cliff of unintended consequences and carry trusting institutions over the edge.

Democracy does not need a machine savior. It needs better light.

To understand the relationship between conscience and tools, let's quickly consider three American luminaries. Each acted from conscience. And in each case, a tool extended human reach:

- Abraham Lincoln's speeches traveled through newspapers and telegraph networks.
- Harriet Beecher Stowe's mass-produced novel, *Uncle Tom's Cabin*, forced readers to see the human reality of slavery.
- Booker T. Washington helped to inspire the Rosenwald Schools Movement, which produced nearly 5000 schoolhouses that expanded educational opportunity for African American children across the rural South.

In each case, human conscience moved first. The tool followed and amplified the human intervention. The tool extended reach, but was never given authority.

That same asymmetry should guide Civic AI today. Humans must pursue civic ends. Systems must be designed so they cannot quietly undermine those ends. Human beings honor democratic principles through intentional action. Civic AI honors democratic principles through designed constraint.

Both may align toward the same civic horizon - but from very different roles. The *Fifty-Gates Democratic Permissions Threshold* instrument makes those boundaries explicit. It treats any Civic AI as a proposed public institution that must earn permission to exist near public power. It insists on non-agency, human moral primacy, reversibility, and anti-weaponization, and it protects pluralism and expressions of conscience. Fail any one of the fifty democratic gates and the AI is disqualified from public service.

The complete *52-Gates* instrument can be accessed at Google Books. I invite all of you to critique, adapt, and extend these ideas in your own work. The *Six Rings* concisely map how digital developments are overwhelming Americans - both individually and across civic life.

All of us - academics, policymakers, AI developers, and citizens - must recognize that agentic AI operating near governance poses a direct threat to 2500 years of democratic evolution.

But a different path is available: non-agentic, democratically-certified Civic AI.

My *Civic Digital Visibility First* principle, together with the *Fifty-Gates Democratic Permissions Threshold*, attempts to reframe the alignment challenge and to systematically test every new AI system before it approaches public power.

Yet civic digital invisibility will continue to remain a structural threat to us all - whether we choose to recognize it or not.

We are at a critical juncture in the emerging field of Civic AI. Let us redefine "Civic AI" clearly - while we still can. Any system that approaches public power without first treating democratic principles as non-negotiable limits is not Civic AI at all; it is simply powerful AI operating near democracy. Civic AI, by definition, begins with constitutional constraint: non-agency, human moral primacy, and refusal-first admissibility. Anything less may be technologically impressive, but it does not qualify for a place near public authority. In other words, *Civic Digital Visibility First* is a governing principle of Civic AI - not a subtype of it.

In the years ahead, human beings - our future Lincolns, and Stowes and Washingtons - will continue to paint the civic canvas. But the role of Civic AI is not to hold the brush.

Edith Wharton once wrote: "There are two ways of spreading light: to be the candle, or the mirror that reflects it."

Let human conscience remain the candle. Then Civic AI may serve as the mirror that reflects that light into the civic world.

Chapter 7: Our Time to Choose

- **"US Civic AI Trial of the Century: Democratic Governance at Risk - Will Civic or Agentic AI Prevail?" (Speculative Courtroom Drama Portraying Real-World Stakes)**

In 2026, as governments are quietly and rapidly integrating autonomous (agentic) AI systems into federal, state, and local offices across the United States, it is tempting to treat these deployments as neutral modernization. This "*U.S. Civic AI Trial of the Century*" is a fictional 2030 courtroom drama designed to challenge that assumption, inviting readers to see today's AI tools not as mere efficiency upgrades, but as constitutional choices about where democratic judgment resides - and what happens if practical responsibility drifts out of human hands. Complacency now may be putting our future at risk. As Abraham Lincoln reminds us: "The dogmas of the quiet past are inadequate to the stormy present… As our case is new, so we must think anew, and act anew."

Literary Introduction of the Creative Analytical Piece,

"U.S. Civic AI Trial of the Century"

Introduction

Christopher Hunt Robertson's *U.S. Civic AI Trial of the Century: Democratic Governance at Risk* operates at the intersection of political philosophy, constitutional law, and speculative fiction. Framed as a fictional federal court case set in the year 2030, this analytical creative piece investigates a pivotal dilemma: "Will Civic or Agentic AI Prevail?" By examining the work through a literary lens, one can discern how Robertson utilizes the conventions of speculative legal drama to construct a cautionary tale about technological acceleration and a rhetorical framework for safeguarding human agency. The text ultimately argues that embedding autonomous systems within public governance threatens to hollow out democratic institutions, proposing instead a restorative "visibility infrastructure" that keeps decision-making power firmly in human hands.

Genre, Setting, and Temporal Framing

The choice of genre-speculative legal fiction - is critical to the work's rhetorical efficacy. By setting the narrative in a 2030 federal courtroom, Robertson establishes a crucible where abstract philosophical principles must confront concrete legal constraints. The courtroom setting naturally demands evidence, adversarial arguments, and structured deliberation, mimicking the very democratic processes that the text seeks to protect.

Furthermore, the temporal positioning of the trial in 2030 provides a necessary distance for defamiliarization, allowing contemporary readers to view current technological trends with greater objectivity. It presents a future that is close enough to feel urgent, yet sufficiently removed to allow for heightened dramatic stakes. This framing transforms a technical policy debate into a high-stakes constitutional crisis, treating the evolution of artificial intelligence not merely as a matter of engineering, but as a core challenge to the social contract itself.

Central Conflict: Agentic AI Versus Civic AI

At the heart of the narrative lies a binary opposition between "Agentic AI" and "Civic AI." Robertson uses this dichotomy to illustrate the psychological and systemic shifts occurring in automated governance.

- **Agentic AI as the Antagonist:** In the text's framework, Agentic AI represents the drive toward autonomous optimization, automated policy recommendation, and citizen scoring. Mechanistically, it functions as a force that stealthily strips citizens and elected officials of their deliberative responsibilities. The narrative portrays this not as an overt algorithmic coup, but as a silent, comfortable migration of authority from human-led public institutions to opaque, optimizing algorithms.

- **Civic AI as the Restorative Counter-Model:** Conversely, Robertson introduces "***Non-Agentic Civic AI***" as the narrative's ideological resolution. Rather than making choices or executing public functions, this model functions strictly as a "visibility infrastructure." It is designed to preserve public records, map system-level pressures, and display aggregate patterns without imposing algorithmic judgments. Through this contrast, the text shifts the definition of "AI alignment" away from technical safety protocols and toward institutional, democratic subordination.

Symbolic Architecture: The Observatory and the Code

Robertson anchors his theoretical concepts in two primary symbolic constructs: the *Franklin – Dewey - Citizenry Civic Data Observatory* and the *Six Rings Zero Acceleration Code*. These elements serve both as narrative world-building devices and as philosophical allusions.

The naming of the *Franklin – Dewey - Citizenry Civic Data Observatory* carries deep historical and philosophical resonance. By invoking Benjamin Franklin, the text recalls Enlightenment ideals of civic virtue, practical invention, and transparent public square governance. The inclusion of John Dewey alludes to pragmatist educational and political philosophy, which emphasizes that democracy is not merely a set of institutional gears but an active, deliberative way of life led by an informed public. The Observatory symbolizes a rights-constrained, transparent, and revocable oversight architecture that restores empirical visibility to the public sphere.

The *Six Rings Zero Acceleration Code* serves as a technocratic Hippocratic Oath for the public sector. The "Zero Acceleration" mandate operates as a vital metaphor for restraint in an era defined by uncritical technological speed. By formalizing these concepts into a legal and ethical code within the trial, the narrative demonstrates how democratic societies can utilize structural roadblocks to assert sovereignty over automated systems rather than passively adapting to them.

Rhetorical Strategy and Tone

The tone of *U.S. Civic AI Trial of the Century* balances analytical gravity with imaginative world-building. It adopts the formal register of legal and academic discourse to establish authority, yet its underlying structure remains deeply creative. The text avoids a simplistic anti-technology stance; instead, it uses a precise, matter-of-fact vocabulary to dissect system design. By presenting the arguments within a trial format, Robertson ensures that both the optimization arguments of autonomous systems and the humanistic defenses of civic infrastructure receive structured articulation, ultimately guiding the reader toward a synthesis that prioritizes civic duty over algorithmic efficiency.

Conclusion

This thought experiment uses speculative literature to diagnose a profound civic vulnerability: the temptation to trade the friction of democratic deliberation for the smooth optimization of autonomous systems. By dramatizing this tension in a fictional 2030 courtroom, *U.S. Civic AI Trial of the Century* underscores that the ultimate risk of artificial intelligence is the erosion of humanity's obligation to govern itself. Through the conceptual models of visibility infrastructure and rights-constrained observatories, the work moves beyond mere critique, offering a structured literary blueprint for a future where technology amplifies, rather than replaces, the democratic public sphere.

US Civic AI Trial of the Century: Democratic Governance at Risk - Will Civic or Agentic AI Prevail? -

Christopher Hunt Robertson, M.Ed.

(With Collaborative Support from ChatGPT, Perplexity, and Gemini/Google)

In the middle decades of the twentieth century, as the United States struggled through the constitutional tensions of the New Deal, Americans did not learn about those debates solely by reading judicial opinions or academic treatises. They learned about them through careful reporting - columns and dispatches that treated constitutional controversy not as an abstract legal exercise, but as part of the nation's civic life.

Arthur Krock, whose *"In the Nation"* column appeared for decades in *The New York Times*, became one of the most respected practitioners of that tradition. He explained difficult constitutional and institutional questions in plain language, writing with a calm analytical voice and a notable preference for understatement over spectacle. His work helped generations of readers understand not only what their government was doing, but what kind of republic it was in the process of becoming.

The pages that follow present a constitutional thought experiment written in that tradition for our age with artificial intelligence. They present an imagined U.S. District Court case, heard in 2030, asking whether democratic judgment in the United States may constitutionally be delegated to advanced Agentic AI systems operating near public authority and, if not, what form of ***Non-Agentic Civic AI*** infrastructure might instead prove admissible under the Constitution. The case itself is fictional. The constitutional questions it explores are not.

This piece invites readers to experience the debate through constitutional journalism. Each report is written as a dispatch from the Court, filed in the measured voice of a correspondent who treats artificial intelligence not as spectacle, but as another chapter in the continuing history of constitutional government and democratic institutions.

Two voices appear throughout these reports.

The first belongs to **Lex Owley**, Chief Congressional Correspondent for the fictional *U.S. Daily Beacon*. His reporting records what happens: who speaks, what arguments are made, what questions the judges ask, and how the proceedings unfold. His voice is human, historical, and journalistic.

The second belongs to a ***Civic Digital Archivist AI***, presented as a ***Non-Agentic Civic AI*** prototype under the proposed stewardship of the Library of Congress. The Archivist neither advises the Court nor predicts outcomes. It neither recommends policy nor evaluates competing constitutional arguments. Its role is narrower, and intentionally so.

Whenever a report introduces a constitutional concept that readers may not immediately recognize - *"civic ceiling," "delegated discretion," "visibility infrastructure," "Non-Agentic Civic AI," "structural drift,"* and related terms—the *Archivist* provides a brief "*Archivist AI Note*" in indented italics. Each note answers a single question: what that term means within the constitutional framework presently under discussion. The notes define, clarify, and preserve the public record, but they never recommend. In this way, the page itself becomes a small demonstration of the civic architecture this volume proposes.

The correspondent reports. The *Archivist* illuminates. The Court interprets. Citizens remain responsible for their own conclusions.

The six news dispatches below serve two purposes.

First, they provide a narrative path for our two primary constitutional questions:

> **Can a democratic republic remain fully self-governing as increasingly capable agentic systems begin exercising practical judgment near public authority?**
>
> **Could *Non-Agentic Civic AI,* designed as civic visibility infrastructure rather than governing intelligence, offers a constitutionally distinct alternative?**

Second, these dispatches prepare readers for a later unveiling of the proposed ***U.S. Non-Agentic Civic AI System***. The Court's imagined opinion is rendered at the conclusion of this series.

Readers should approach the reports that follow not as predictions, but as an exercise in constitutional imagination - an opportunity to consider how the questions now emerging around artificial intelligence might appear if they were ever argued before the nation's high courts.

If the experiment succeeds, the fictional framework will gradually fade from view. What will remain are the constitutional questions themselves - and the responsibility of democratic citizens to answer them.

U.S. DAILY BEACON - Special Constitutional Report One - Monday, August 3, 2030

By Lex Owley, Chief Congressional Correspondent

WHY IS THE U.S. FEDERAL DISTRICT COURT HEARING THIS CASE?

WASHINGTON, D.C. - Civic artificial intelligence, and its growing role in the ordinary administration of government, has steadily drawn the attention of citizens, lawmakers, and public institutions across the country. That attention intensified this week when challengers to the federal government's expanding use of agentic artificial intelligence secured an unusual three-judge district court to hear what many observers regard as one of the most significant constitutional technology cases ever brought before a federal court.

The plaintiffs contend that the increasing delegation of governmental authority to Agentic AI systems threatens the fair administration of public accommodations, employment, education, federally funded programs, voting rights, public benefits, and access to public facilities throughout the United States. They argue that these questions extend well beyond technology policy and reach the constitutional structure of democratic government itself.

To the surprise of many legal observers, Congress recently restored, in carefully limited form, a rarely used three-judge district court procedure once associated with major Civil Rights Era cases and later employed only sporadically in a handful of nationally significant civil rights actions. Concluding that constitutional challenges involving the nationwide use of advanced agentic artificial intelligence required an equally deliberate form of judicial review, lawmakers enacted the Civic AI Constitutional Review Act of 2029. The law directs that such cases be heard by a three-judge federal court in Washington, D.C., with any appeal proceeding directly to the Supreme Court.

Long before the Court convened this morning, the public line extended beyond the marble plaza and around the east side of the building. College students stood beside retired factory workers. Software engineers chatted quietly with constitutional historians. Young parents arrived pushing strollers while several veterans unfolded lawn chairs along the sidewalk. Business attire mixed easily with work boots, backpacks, and summer dresses. Few appeared to know one another. All had come for the same hearing.

As today's proceedings began in what some have already called the "Civic AI Trial of the Century," one question echoed through both the press gallery and the crowd outside the Court:

Why this case - and why now?

The complaint before the Court does not ask whether artificial intelligence can write code, answer questions, or manage routine office tasks. It asks something narrower, and considerably more difficult: whether democratic government may constitutionally delegate consequential governmental judgment to agentic artificial systems while preserving the conditions of self-government the Constitution was designed to protect.

In accepting jurisdiction, the Court agreed to consider whether a constitutional boundary exists between machine-assisted civic analysis and delegated machine judgment exercised near public authority - a boundary some scholars have begun referring to as the "civic ceiling."

*Archivist AI Note: Within the constitutional framework presented to the Court, the **civic ceiling** denotes the boundary between machine-assisted civic visibility and delegated governmental judgment. Systems operating below this boundary illuminate public information without exercising consequential public discretion. Systems operating above it begin exercising judgment in domains the petitioners argue remain constitutionally non-delegable. This note records terminology used during the proceedings and expresses no view regarding its legal merits.*

In that sense, today's arguments open another chapter in the continuing history of constitutional government in the United States.

From Technical Debate to Constitutional Question

For several years, technical experts, administrators, and vendors have described advanced artificial intelligence in government primarily in terms of efficiency and modernization. Systems that route cases, rank risks, score eligibility, and recommend policies have been presented as tools to reduce backlogs, standardize decisions, and improve administrative responsiveness.

The case now before the Court presents a different constitutional question. It argues that once artificial systems begin exercising delegated discretion in matters involving rights, eligibility, enforcement, and access to public benefits, they no longer function merely as administrative tools. They begin altering where public judgment occurs.

According to the filings, democracies must distinguish between systems that help citizens and officials see public reality more clearly and systems that increasingly substitute for the human act of judging itself.

*Archivist AI Note: Within this framework, **Non-Agentic Civic AI** functions as civic visibility infrastructure. It organizes, preserves, and illuminates publicly available information without exercising governmental judgment. Whether such a distinction possesses constitutional significance is one of the questions now before the Court.*

The distinction now appears prominently throughout the arguments the Court has agreed to hear.

Echoes from the New Deal Era

Observers have frequently compared the present moment to the constitutional struggles surrounding President Franklin D. Roosevelt's New Deal, when the nation debated whether new forms of economic and administrative infrastructure could be reconciled with existing constitutional limits. Then, as now, the central question was not technology itself, but power.

New Deal programs raised questions about the scale of federal intervention, the delegation of authority to new agencies, and the durability of constitutional limits once new mechanisms of decision were put in place. Journalists such as Arthur Krock helped readers understand not only what government was doing, but what kind of republic it was becoming.

Today's case arrives in a similar spirit. The filings ask whether a constitutional republic can safely incorporate increasingly capable Agentic AI systems into the machinery of governance without gradually altering who exercises practical public judgment and how citizens may understand, question, and contest its use.

One retired historian waiting outside the Court summarized the concern in plain terms: "We have always known that tools change what institutions can do. This case asks whether new tools might quietly change what those institutions become."

An Observatory in the Gallery

Among those quietly observing today's proceedings was a small Library of Congress delegation planning the proposed ***Franklin - Dewey - Citizenry Civic Data Observatory***, attending not as participants in the litigation, but as the federal government's primary collector, preserver and curator of national memory. The delegation included a ***Civic Digital Archivist AI***, a prototype responsible solely for preserving the historical record of the hearings, and two members of the ***Civic Digital Telescopic AI Ensemble*** - also original models in testing mode.

Project officials emphasized that these purpose-built AI systems possess no authority to advise the Court, recommend outcomes, or influence the proceedings in any manner. Their purpose, according to project materials, is to organize and illuminate publicly available information while remaining under strict non-agentic constitutional constraints.

When approached by reporters regarding the pending litigation, an Observatory representative offered a brief statement: "No comment at this time. Descriptive analytical reports, without recommendations, will be released publicly following completion of the proceedings in accordance with Library of Congress publication guidelines."

Several reporters smiled before returning to their notebooks.

What Is at Stake

Although the docket lists specific statutes, programs, and governmental deployments now under challenge, the constitutional questions extend well beyond any single agency or technological platform.

Among them are whether democratic governments may constitutionally rely upon systems exercising delegated discretion beyond meaningful public oversight; whether public authority may, in practice, rest upon models that no elected official fully understands; and whether a constitutional republic can preserve meaningful human responsibility if practical judgment gradually migrates into machine behavior.

Supporters of Agentic AI systems argue that such tools have become necessary to manage the scale and complexity of modern government. Critics respond that constitutional democracies have historically paid a high price whenever significant public authority has drifted beyond effective human supervision.

As spectators entered the courtroom, several glanced almost simultaneously at their phones, smiling or scowling as they scrolled or texted.

Moments later, the clerk's announcement brought the Court to order. Conversation ceased as every eye turned toward the bench.

A Trial About Responsibility

Outside the courtroom, a law student offered a succinct perspective: "This is not a case about whether machines can help us see more clearly. We already know they can. It is a case about whether we intend to keep moral responsibility where the Constitution placed it, in human hands that can answer to the public."

No major public figure stands accused in these proceedings, and no scandal dominates the record. Instead, the Court has been asked to consider questions concerning the future architecture of democratic governance and the location of constitutional responsibility in an age of increasingly capable digital systems.

Looking Ahead

The three-judge panel is expected to hear arguments from government representatives, civil liberties advocates, technical experts, and historians.

In the days ahead, this paper will report the proceedings as they unfold, recording the constitutional questions placed before the Court and the arguments offered in response.

For now, the lines remain long, the questions remain open, and the Court's work has only begun.

U.S. DAILY BEACON - Special Constitutional Report Two - Tuesday, August 4, 2030

By Lex Owley, Chief Congressional Correspondent

CAN DEMOCRATIC JUDGMENT BE DELEGATED TO AI?

WASHINGTON, D.C. - If yesterday's opening hearing established the size of the controversy, today's proceedings began to clarify its central constitutional question.

Can a democratic republic permit artificial systems to participate in governmental judgment without gradually transferring into those systems responsibilities the Constitution expects accountable human officials and institutions to retain?

That question surfaced repeatedly throughout today's arguments, though often in more technical language. Counsel referred to delegated discretion, Agentic AI systems, constitutional limits, and what one scholar observing the proceedings has begun calling the *civic ceiling*.

Archivist AI Note: Within the constitutional framework presented to the Court, ***delegated discretion*** *refers to authority transferred from accountable public officials to systems capable of selecting among consequential alternatives within a defined scope. The concern is not whether such systems are technically capable, but whether delegated operational judgment remains compatible with democratic accountability once those systems begin shaping public outcomes. This note records terminology used during the proceedings and expresses no view regarding its legal merits.*

The morning's exchange suggested that the Court is not being asked whether computers may assist public officials. That question has largely been settled in practice. Governments routinely rely upon databases, communications networks, records systems, scheduling software, and many other digital tools that organize information and support routine administration.

The more difficult constitutional question is whether assistance remains assistance once systems begin ranking, filtering, routing, prioritizing, or otherwise structuring the practical field within which public judgment occurs.

One judge asked whether the line proposed by the challengers was merely philosophical. Another asked whether it could be translated into an administrable rule of constitutional law. A third returned repeatedly to a simpler question: "When does support become substitution?"

Archivist AI Note: Within the present framework, the ***civic ceiling*** *denotes the boundary separating machine-assisted civic visibility from delegated governmental judgment. Systems operating below that boundary illuminate public conditions without exercising consequential public discretion. Systems operating above it begin exercising practical judgment in domains such as rights, eligibility, enforcement exposure, agenda-setting, or access to public benefits. This note is descriptive only and expresses no opinion regarding the constitutional validity of the concept.*

Institutional Drift

Counsel for the petitioners argued that the Constitution concerns itself not only with whether governmental decisions can later be explained. It also considers how those decisions were reached, who remained responsible for them, and whether citizens could meaningfully understand and contest the chain of public judgment.

That argument attracted attention because it did not depend upon dramatic failures or speculative future disasters. Instead, counsel described a gradual institutional progression: A system is introduced to help sort. Then to recommend. Then to route. Then to flag exceptions. Then to identify priority cases.

Eventually, public officials find themselves working within categories and sequences the system has already arranged.

The formal structure of government remains unchanged, but the practical location of judgment quietly begins to migrate.

Several members of the Court appeared interested in whether such institutional drift represented ordinary administrative convenience, or whether it gradually alters the constitutional character of governance itself.

*Archivist AI Note: Within these proceedings, an **agentic system** is understood as one capable of selecting actions, advancing institutional processes, or exercising delegated operational discretion without fresh human judgment at every consequential step. The defining issue is delegated public discretion, not computational sophistication.*

Public Intelligibility

The discussion soon turned toward a second constitutional concern.

One judge asked whether democratic legitimacy requires governmental reasoning to remain publicly intelligible even as institutions become increasingly complex.

Counsel replied that constitutional accountability depends not merely upon the existence of human supervisors, but upon citizens' continuing ability to understand who exercised judgment, on what basis, and under what authority.

*Archivist AI Note: Within this framework, **public intelligibility** refers to the ability of citizens, courts, legislators, and public officials to understand the institutional path by which consequential governmental decisions are reached. It concerns democratic contestability rather than administrative convenience.*

Exploring Constitutional Alternatives

One notable development today was the degree to which members of the Court began asking not only what constitutional limits might require, but what constitutional alternatives might permit.

If certain forms of delegated machine judgment prove constitutionally suspect, what remains available to governments confronting problems that increasingly unfold at digital speed?

That question allowed counsel to distinguish more clearly between governance systems and visibility systems.

Nothing in the petition, counsel argued, would prevent governments from employing advanced computational systems to preserve records, detect aggregate patterns, illuminate emerging conditions, or improve shared civic understanding.

The proposed constitutional line does not forbid machine-assisted visibility. It limits delegated machine judgment near public authority.

That distinction appeared to interest some judges because it offered a path beyond the increasingly common assumption that society must choose either unrestricted adoption or wholesale rejection of artificial intelligence.

The constitutional question may instead concern different constitutional designs for incorporating technology into democratic life.

Testing the Line Between Seeing and Deciding

By late afternoon, the atmosphere within the courtroom had changed in a subtle but noticeable way. On the opening day, the case seemed unusually large. Today, it began to sound more recognizably judicial.

The judges were no longer asking whether artificial intelligence would become part of American life. They were asking what constitutional principles must continue governing public authority once increasingly capable digital systems move nearer to it.

That is slower work than public controversy often prefers. It is also precisely the kind of work constitutional courts exist to perform.

Outside the building, several spectators expressed frustration that the Court was not yet moving toward simple, clear distinctions. Inside, the rhythm of questioning suggested that the judges were carefully examining the case's layered complexities. They appeared to be asking whether a constitutional boundary exists, what values it protects, and how that boundary might be recognized before gradual institutional change makes the question more difficult to answer.

Why Today's Questions Matter

For citizens following the proceedings, today's arguments offered one important clarification. The Court is no longer asking only whether artificial intelligence can enter government.

It is asking what becomes of responsibility, explanation, and democratic contestability once artificial systems begin shaping the practical field in which governmental judgment occurs. That is a subtler question than much public discussion has recognized.

It may also prove the more enduring one.

If constitutional democracies lose the ability to identify who judged, who decided, and upon what visible basis those judgments were made, formal institutions may remain unchanged even while practical responsibility quietly migrates elsewhere.

The Court did not indicate today whether that migration has already crossed a constitutional line. It asked enough questions, however, to suggest that at least some members regard the issue as one requiring careful judicial examination.

Tomorrow's hearing is expected to focus on the government's affirmative case for Agentic AI systems - the constitutional responsibilities of effective administration, national security, operational scale, and the practical burdens of governing a modern republic.

For now, today's proceedings leave one question more clearly before the nation than yesterday. The issue may not ultimately be whether Americans desire more intelligent machines. It may be whether democratic judgment itself can be delegated without gradually changing the constitutional character of democratic government.

U.S. DAILY BEACON - Special Constitutional Report Three - Wednesday, August 5, 2030

By Lex Owley, Chief Congressional Correspondent

WHAT THE JUDGING PANEL ASKED TODAY

WASHINGTON, D.C. - Courtroom spectators settled quickly into their seats this morning. Attorneys arrived carrying noticeably thicker binders than during the opening days of argument, while several members of the public gallery quietly compared notes from yesterday's proceedings. When the judges entered, conversation ceased almost immediately.

By the third day of hearings in the Court's Civic AI case, the broad constitutional dispute had narrowed into a practical judicial inquiry.

If there is to be a constitutional boundary between machine assistance and machine governance, how is a court to recognize it?

That question, in one form or another, guided much of today's argument.

The judges spent little time asking whether artificial intelligence is powerful, useful, or inevitable. They asked instead whether the distinctions proposed by the challengers could be expressed clearly enough to become administrable constitutional law.

Several members of the Court repeatedly returned to one underlying concern. Can a constitutional republic preserve meaningful human responsibility if its institutions increasingly rely upon systems that shape practical outcomes before any accountable public official makes the final visible decision?

Archivist AI Note: Within the present constitutional framework, ***final visible decision*** *refers to the identifiable point at which accountable human authority continues to bear responsibility for a consequential governmental act. The concern raised during today's proceedings is whether earlier stages of ranking, filtering, scoring, or routing may so structure available choices that the remaining human decision becomes largely formal. This definition is recorded solely to clarify terminology used before the Court.*

Drawing the Constitutional Line

As in many constitutional cases, today's arguments focused less on broad agreement than on line-drawing.

One judge asked whether the challengers' framework would prohibit only systems making final governmental decisions, or also systems that substantially influence those decisions before they reach a human official.

Another questioned whether ranking itself should be considered delegated judgment if ranked recommendations consistently determine which matters receive immediate attention and which quietly recede into the administrative background.

A third judge proposed a hypothetical system that never issues commands, never signs an order, and never formally denies a benefit, yet routinely sorts applications into urgency tiers that busy caseworkers almost always follow.

Would such a system constitute assistance, or governance by another name?

The question was put quietly. It seemed to sharpen the attention of everyone present.

Counsel for the challengers replied that constitutional analysis cannot concern itself only with ceremonial endpoints. If a system structures practical outcomes involving rights, eligibility, enforcement, or civic participation, they argued, the location of meaningful judgment may already have shifted, even if a human signature remains on the final page.

Archivist AI Note: Within the framework presented today, ***meaningful judgment*** *refers not merely to nominal human review but to the continuing capacity of accountable officials to examine, depart from, contest, and assume responsibility for consequential governmental determinations without being materially constrained by pre-structured system outputs. This definition records terminology used in today's proceedings and expresses no view regarding its constitutional merits.*

The Problem of Administrable Standards

If the challengers encountered their most sustained questioning today, it came when several judges pressed them to describe a constitutional standard that lower courts and public agencies could consistently apply.

Several attorneys exchanged handwritten notes before counsel returned to the lectern.

How, one judge asked, is a court to determine when influence becomes delegation? Is the constitutional line crossed at automatic execution? At ranking? At recommended action? At institutional dependence? At opacity? Or only when a system becomes practically impossible to override?

The challengers resisted the suggestion that one mechanical trigger could resolve every future case. Instead, counsel proposed a series of constitutional indicators: whether a system crosses the civic ceiling; whether it operates within domains of consequential public discretion; whether officials predictably defer to its outputs; and whether citizens retain meaningful opportunities to identify, challenge, and reverse the judgment affecting them.

That response did not resolve every question from the bench. It did, however, move the discussion toward familiar constitutional territory. Courts have long examined questions of structure, accumulated effect, and institutional function without relying upon a single bright-line test. Several courtroom observers later remarked that today's questioning appeared to explore whether this case belongs within that tradition.

Can Influence Become Authority?

One question resurfaced repeatedly throughout the afternoon. Must a system issue binding commands before it becomes constitutionally significant?

Government counsel maintained that many AI systems presently under challenge remain advisory. They summarize files, identify anomalies, estimate probabilities, compare cases, and help public officials manage overwhelming volumes of information. If human officials remain free to reject those recommendations, the government argued, constitutional authority remains where it has always resided - with accountable public officers.

A judge then asked whether formal freedom to disagree retains practical meaning within institutions where speed, workload, and organizational expectations make disagreement increasingly uncommon. **If officials follow recommendations in ninety-seven percent of cases, the judge asked, does the system remain merely advisory?**

No immediate answer entirely resolved the question.

Government counsel replied that high rates of agreement may simply reflect the usefulness of the technology. Counsel for the challengers responded that usefulness itself may become the mechanism through which practical discretion gradually migrates away from visible human judgment.

Archivist AI Note: The constitutional issue discussed here is not determined solely by whether a system is formally labeled "advisory." Within the framework presented to the Court, an advisory system may nevertheless become practically authoritative if institutional actors predictably organize their work around its outputs and rarely depart from its recommendations. This definition is recorded solely for public understanding.

That exchange appeared to become one of the day's central inquiries. Rather than distinguishing neatly between human government and machine government, the Court appeared increasingly interested in the constitutional gray zone where one may gradually begin to resemble the other.

The Visibility Question

A second line of questioning focused not only upon delegated judgment, but upon public visibility.

Several judges asked whether ordinary citizens remain able to understand how governmental outcomes are reached once AI-mediated systems assume larger roles in administrative processes.
If a citizen loses access to a benefit, receives heightened scrutiny, or moves lower within an administrative queue, what explanation must government still be able to provide?

One judge questioned whether explanations offered after the fact remain sufficient if no public official can clearly describe how practical outcomes were reached. Another asked whether democratic confidence can endure if agencies themselves become unable to explain, in ordinary language, why one citizen received attention before another.

The challengers argued that democratic legitimacy depends not only upon technical reviewability, but also upon public intelligibility.

*Archivist AI Note: Within today's proceedings, **public intelligibility** refers to the continuing ability of affected citizens, courts, oversight bodies, journalists, and the public to understand what governmental action occurred, by what authority, and through what identifiable process. This definition is recorded solely to clarify terminology presented before the Court.*

Exploring an Alternative

One notable development today was the Court's growing interest not only in what constitutional limits might exist, but in what governmental uses of AI might remain fully permissible.

If certain forms of agentic delegation prove constitutionally problematic, what technological alternatives remain available?

Counsel responded by drawing a clearer distinction between governance systems and visibility systems. Governments, they argued, remain free to preserve records, detect patterns, organize information, illuminate public conditions, and improve institutional awareness without delegating consequential governmental judgment.

*Archivist AI Note: Within the present framework, **visibility infrastructure** refers to systems designed to increase public understanding without exercising consequential governmental judgment. Such systems organize, preserve, detect, and illuminate publicly available information while remaining structurally outside the chain of governmental decision-making. This definition is recorded solely for public understanding.*

Several members of the Court appeared particularly interested in that distinction because it offered a constitutional alternative to the all-or-nothing debate that has often characterized public discussion of artificial intelligence.

Preserving Accountable Human Judgment

By late afternoon, the rhythm of the courtroom had changed. On the first day, the proceedings seemed unusually large. On the second, unusually theoretical. Today, they sounded recognizably judicial.

The judges were no longer asking whether artificial intelligence belongs in the future. They were asking what constitutional standards might preserve accountable human judgment once such systems move near public power.

That is slower work than public controversy often prefers. It is also precisely the kind of work constitutional courts exist to perform.

Outside the building, several spectators compared their own expectations with the pace of the proceedings. One remarked that the Court had not yet produced a simple rule. Another quietly replied that constitutional questions seldom do.

What Citizens Learned Today

For citizens following the proceedings from a distance, today's hearing offered one important clarification. The Court is no longer asking only whether artificial intelligence may enter government.

It is asking what becomes of responsibility, explanation, and public contestability once AI begins shaping the field within which governmental decisions are made. That is a subtler question than much public discussion has recognized, and it may also prove the more enduring one.

If constitutional democracies lose the ability to identify who judged, who decided, and upon what visible basis those judgments were made, formal institutions may remain unchanged even while practical responsibility quietly migrates elsewhere.

The Court did not indicate today whether that migration has already crossed a constitutional line. It asked enough questions, however, to suggest that at least some members regard the issue as one requiring careful judicial examination.

Tomorrow's hearing is expected to focus more directly upon the government's affirmative case for Agentic AI systems: administrative necessity, efficiency, national competitiveness, and governing at digital scale.

For today, one observation seemed sufficient. The Court devoted much of its attention not to whether artificial intelligence should exist, but to where constitutional judgment must remain.

U.S. DAILY BEACON - Special Constitutional Report Four - Thursday, August 6, 2030

By Lex Owley, Chief Congressional Correspondent

GOVERNMENT'S CASE FOR NECESSITY

WASHINGTON, D.C. - Government counsel arrived this morning carrying several additional binders, while courtroom staff placed two demonstration exhibits beside the lectern before proceedings began. By the time the Marshal called the Court to order, it was already apparent that today's arguments would differ noticeably from those presented earlier in the week.

If the first several days of the Civic AI case largely examined the constitutional risks of delegated judgment, today the government was given a fuller opportunity to explain why it believes certain forms of Agentic AI have become practically necessary in modern governance. The result was a more complete picture of the pressures confronting public institutions.

Government counsel did not present advanced AI systems as constitutional ideals, or as constitutionally neutral. They presented them as responses to mounting administrative strain: growing backlogs, emergency conditions, increasingly sophisticated digital threats, and public expectations of speed that many existing governmental structures were never designed to meet.

Their central constitutional argument was straightforward. In a nation of more than 330 million people, with complex benefit systems, dense regulatory frameworks, and continuous digital activity, public institutions now confront volumes of information and rates of change that human staff alone struggle to manage. Without some measure of delegated machine initiative, counsel argued, critical governmental functions may simply fail to perform at the level modern citizens increasingly expect, and the Executive's duty to execute the laws effectively may be placed at risk.

Archivist AI Note: Within today's proceedings, ***administrative strain*** *refers to persistent institutional pressures arising from increasing workload, limited staffing, statutory complexity, emergency conditions, and rising public expectations regarding governmental responsiveness. The term describes conditions discussed before the Court and expresses no evaluative judgment regarding proposed remedies.*

The Court Weighs Administrative Reality

Much of today's presentation relied upon examples familiar to anyone who has spent time waiting for a public agency to respond.

Government counsel described unemployment systems processing millions of claims during periods of economic disruption, public-health agencies receiving dense streams of information during disease outbreaks, social-service offices managing caseloads too large for detailed manual review, and cybersecurity teams monitoring networks where significant events unfold in fractions of a second. In each example, counsel argued, Agentic AI systems were introduced not to seize authority for their own sake, but to prevent institutional failure.

One illustration involved automated triage systems. Machine systems, counsel explained, classify incoming matters into routine, urgent, and critical categories, allowing limited human staff to concentrate first where probable need appears greatest.

Several members of the public gallery quietly nodded as these examples were presented.

If such systems were removed entirely, government counsel argued, backlogs would lengthen, emergencies could be overlooked, and delays already familiar to many citizens would become substantially worse, raising concerns not only about inefficiency, but also about practical harm to due process.

Archivist AI Note: Within today's proceedings, a ***triage system*** *refers to the classification of incoming cases, claims, or signals into priority categories according to predefined rules or learned patterns. The constitutional question before the Court is whether such classification constitutes permissible administrative assistance or delegated practical judgment when it influences access to governmental attention. This definition is recorded solely for public understanding.*

National Security and Cyber Threats

The government's strongest presentation came during discussion of national security and digital infrastructure.

Counsel described computer networks subjected to continuous automated probing by hostile actors operating at machine speed. Under such conditions, counsel argued, human monitoring alone can no longer provide adequate protection.

One judge asked whether such systems differ meaningfully from earlier automatic safeguards, including circuit breakers or intrusion alarms.

Government counsel replied: "The scale and adaptability of these threats have changed in kind, not merely in degree."

The government's position was that defensive systems must sometimes detect anomalous behavior, isolate suspicious activity, and temporarily restrict access before human review can occur. Counsel acknowledged that such systems may operate across the civic ceiling, but argued that carefully constrained defensive exceptions remain necessary if government is to protect essential public infrastructure at all.

Even advocates of ***Non-Agentic Civic AI*** conceded that exceptions would be needed. Without Agentic AI, military, intelligence, and law enforcement agencies would sometimes be unable to carry out their duties. But then it was noted that only about ten percent of all combined federal, state, and local government employees in the U.S. work within such agencies.

The Efficiency Argument

Not all of today's defense rested upon emergencies. Government counsel returned repeatedly to the ordinary work of government: document preparation, eligibility pre-screening, scheduling, routine correspondence, and matching citizens with programs for which they may qualify.

Counsel spoke without rhetorical flourish.

The claim was not that such systems make government wiser. Rather, in many public offices, they make government possible.

One judge asked whether language emphasizing efficiency risks transforming constitutional questions into budgetary ones. If Congress could instead hire additional employees, is automation a constitutional necessity, or an appropriations choice?

Government counsel replied that the distinction often proves difficult in practice. Public expectations, available funding, statutory responsibilities, and technological capabilities evolve together. Agencies, counsel argued, frequently seek automation not because constitutional principles have changed, but because operational demands have - and because the constitutional duty to execute the laws in a timely and effective manner cannot easily be met with existing human capacity alone.

Competitiveness in a Changing World

Government counsel also raised a broader concern. Other nations, including governments operating under markedly different constitutional traditions, continue expanding the use of advanced Agentic AI systems in administration, infrastructure, and national security.

A judge asked whether constitutional interpretation should be influenced by international technological competition.

Government counsel responded that while the Constitution remains controlling, constitutional government must also remain capable of functioning within an increasingly complex and rapidly changing world. A republic that refuses tools necessary to manage modern conditions, counsel suggested, may risk falling short of its own obligations to protect the public and maintain effective governance.

The point was raised briefly, but remained quietly present throughout the afternoon.

The Human Element

One of today's memorable moments concerned not technology, but the people who work inside government.

Government counsel described public employees managing large caseloads, emotionally demanding responsibilities, and constant public scrutiny under increasingly constrained resources. Several agency representatives seated behind counsel quietly exchanged glances as the discussion turned toward staffing pressures and burnout.

A judge asked whether these realities should matter in constitutional analysis. Government counsel answered that they should. A constitutional republic, counsel argued, must care not only about institutional limits, but also about the practical conditions under which human beings are asked to uphold those limits each day. When employees become chronically overwhelmed, the risks include not only inefficiency, but error, exhaustion, and gradual withdrawal from responsibility itself.

Within that context, counsel argued, carefully designed Agentic AI tools may strengthen rather than diminish human responsibility by reducing burdens that need not remain exclusively human, allowing officials to focus attention on decisions that must remain fully accountable and visible.

Archivist AI Note: Today's proceedings briefly addressed the emotional and cognitive burdens experienced by public servants operating under sustained workload and time pressures. Within the proposed Non-Agentic Civic AI framework, such burdens remain inherently human responsibilities. Visibility-oriented Civic AI is intentionally constrained so that it informs human judgment without attempting to internalize or replace the moral responsibilities carried by public officials. This note records the distinction without expressing any constitutional conclusion.

A Fair Hearing of Necessity

For citizens following the proceedings from outside the courtroom, today's arguments offered an important perspective. Government's interest in Agentic AI was presented not as technological enthusiasm, but as an institutional response to practical pressures that many citizens already recognize: growing workloads, digital threats, rising expectations, and budgets that do not always keep pace with public demand. Within that environment, systems capable of acting with limited initiative may appear less as technological luxuries than as operational necessities.

The Court endorsed none of those arguments today.

It did, however, allow them to receive a full constitutional hearing.

Balancing Capacity and Accountability

By the close of today's session, the overall structure of the case had become increasingly clear. The opening hearings identified the constitutional questions before it. Subsequent arguments explored delegated judgment, structural drift, and the challenge of drawing workable constitutional boundaries. Today's proceedings allowed the government to explain why it believes certain forms of agentic capability have become operationally necessary within modern public administration.

The case has gradually come to resemble less a debate over software than a constitutional tension between the demands of efficient execution and the demands of visible, accountable judgment.

Tomorrow's arguments are expected to turn toward a different question. If constitutional limits exist on delegated governmental judgment, what practical alternative remains?

Challengers supporting proposed ***Non-Agentic Civic AI*** contend that visibility-centered civic infrastructure can relieve many institutional pressures without crossing the boundary between illumination and judgment.

Government counsel has suggested that visibility alone may prove insufficient in certain domains.

The Court has yet to hear that question fully argued.

For today, one observation seemed sufficient. Both sides described themselves as seeking to preserve responsible democratic government. Their disagreement concerns the institutional means by which that responsibility is best sustained.

U.S. DAILY BEACON - Special Constitutional Report Five - Friday, August 7, 2030

By Lex Owley, Chief Congressional Correspondent

THE ALTERNATIVE APPEARING IN THE RECORD

WASHINGTON, D.C. - After several days examining the risks and claimed necessities of Agentic AI in government, today's hearing shifted focus.

The Court began exploring whether a constitutional alternative exists: a class of civic systems designed not to exercise judgment, but to restore visibility.

In plain terms, the petitioners' case no longer rests on criticism alone. It now rests on an assertion that the country possesses a viable non-agentic path for using advanced digital tools near power without allowing those tools to become governing agents.

That alternative has appeared in the record under several names: "*Non-Agentic Civic AI*," "*visibility infrastructure*," "*civic digital lighthouses*," and the proposed "*Franklin – Dewey - Citizenry Civic Data Observatory*." Today, those terms moved from theory into practical description.

A Lighthouse, Not a Helm

Counsel for the challengers began by repeating a simple metaphor that has now surfaced in several filings.

Democracies, they argued, do not need machine helms. They need better lighthouses.

The distinction was more than rhetorical. It described the proposed function of Civic AI in constitutional terms: systems designed to see, not to steer.

Under questioning from the bench, counsel emphasized three features of the alternative they are asking the Court to recognize. Civic AI would be structurally non-agentic. It would operate as visibility infrastructure, not as decision infrastructure. And it would preserve all consequential judgment, authority, and responsibility as fully human.

One judge asked whether this is merely a restatement of ordinary data systems. Counsel replied that it is not. The proposal, they said, is for a dedicated civic data observatory - a public institution designed to maintain continuous, shared visibility into complex digital conditions while explicitly refusing every design path that would allow it to judge citizens, rank persons, or act on its own initiative.

*Archivist AI Note: "**Non-Agentic Civic AI**" is defined, in the framework advanced here, as a strictly non-agentic class of public visibility infrastructure that extends human perception of civic conditions through counts, timestamps, distributions, correlations, and mappings while remaining structurally incapable of interpreting, recommending, deciding, or acting. Its field of view is aggregate and system-level; it is architecturally blind to the targeting, scoring, or surveillance of individual persons. This concept is noted solely to clarify its use in these proceedings.*

A Judge Asks for a Simpler Architecture

As the hearing progressed, one member of the Court posed a more practical question.

"Counsel," the judge said, "you've described visibility infrastructure, delegated judgment, and constitutional boundaries. I'm still trying to understand how these pieces relate institutionally. Is there a simpler way to describe the architecture you're proposing?"

Counsel paused briefly before replying.

"Yes, Your Honor. Imagine a civic lighthouse enclosed within protective walls. The lighthouse extends visibility beyond ordinary sight. The walls protect both the lighthouse itself and the constitutional order it serves - not only against sudden assault, but also against the slow erosion that can gradually weaken democratic foundations."

"When government requires bounded agentic capabilities, accountable human officials may pass through the civic gates to engage external systems for specific clerical or defensive tasks. When those tasks conclude, human authority terminates the engagement and returns within the protected civic structure. Agentic AI may assist the constitutional order from outside, but may never enter the walls that protect the illuminating lighthouse or the human locus of constitutional judgment inside."

Several members of the public gallery later remarked that the analogy appeared to simplify a distinction that had occupied much of the week's testimony.

The Observatory Takes Shape

Within that framework, counsel described the proposed ***Franklin – Dewey - Citizenry Civic Data Observatory*** in careful, practical terms. It would not sit above government. It would sit alongside it, inside the gated civic sphere, observing, recording, and analyzing.

Its tasks would include:

- Continuously assembling lawful, civic-relevant digital signals across domains: health, infrastructure, communication, economic stress, and other indicators.
- Presenting those signals in structured, non-interpretive forms: counts, timelines, distributions, and pattern maps.
- Surfacing emerging pressures, drifts, and anomalies that might otherwise remain invisible to institutions operating at human speed.
- Preserving time-stamped records of civic conditions for later human review.

Whatever else it might become, counsel insisted, it would not become a governing mind.

A judge asked whether this is simply another word for a "dashboard." Counsel replied that the term is more precise than that. The Observatory would be governed by strict constitutional operating constraints: no recommendations, no rankings of individuals, no automatic triggers of enforcement, and no optimization toward any agenda. Its purpose would be to make the civic field visible, not to decide how that field should be changed.

*Archivist AI Note: The phrase "**civic digital lighthouse**" has been used in the filings to describe Civic AI systems whose sole vocation is to illuminate hazards, pressures, and patterns in the civic digital field without intervening in those conditions. In this metaphor, light is treated as a prerequisite for safe navigation, not as a steering mechanism. The Observatory records this metaphor here as part of the descriptive vocabulary in the case.*

Bounded Agentic Tools Outside the Walls

Under further questioning, counsel emphasized that this visibility-first design does not require government to abandon routine automation.

Public institutions, they argued, would remain free to employ bounded Agentic AI tools for clerical execution - document preparation, standardized form completion, scheduling assistance, and other non-discretionary administrative tasks - so long as those tools are initiated, supervised, and concluded by accountable human authority.

Within the walled lighthouse analogy, such systems operate beyond the walls. Human officials pass through the civic gates to engage them for specific missions and return when those missions have been completed. Agentic tools may assist constitutional government in execution; they do not become the locus of constitutional judgment.

The proposed Observatory would remain separate from any agentic system carrying out governmental work. Its reports would be public, available to citizens, agencies, legislators, and the courts alike. But the Observatory would remain structurally separated from outside Agentic AI systems, and would never become part of the chain by which Agentic AI systems carried out governmental tasks.

Human officials alone decide when to engage external systems, what work is assigned beyond the walled perimeter, and when that engagement ends. In that way, Civic AI remains a source of public visibility, while constitutional responsibility continues to rest with accountable human institutions.

Continuous Civic Digital Resolution

The Court appeared particularly interested in how such systems would work over time. Counsel described what they called "***continuous civic digital resolution***" - an ongoing ability to see how conditions in the civic field change, rather than relying only on episodic reports or crises.

In that conception, the Observatory would function less like a camera than a weather service. It would track patterns across days, weeks, and years, revealing where pressures accumulate, where institutional responses cluster, where civic visibility thins, and where digital developments begin to outpace traditional oversight.

One judge asked how such a system would avoid becoming a surveillance instrument.

Counsel answered that the Observatory's field of view is deliberately set at the aggregate and system level. "It would not watch individuals," they said. "It would watch systems." It would see shifts in distributions, not movements of identifiable persons; changes in volumes, not trajectories of single lives.

*Archivist AI Note: The "**aggregate and system-level field of view**" described here reflects a design constraint under which Civic AI observes patterns, flows, and conditions without constructing or maintaining person-level profiles. This constraint is central to the proposed distinction between visibility infrastructure and surveillance infrastructure. No position on the constitutional merits of this distinction is expressed in this note.*

Historical Orientation: Franklin and Dewey

Today's hearing also touched upon the Observatory's historical duties.

Counsel argued that visibility is not only spatial and statistical. It is also temporal.

The proposed Observatory would maintain what they described as "***Time-Triangulated Civic Situational Awareness***" - continuous awareness of past memory, present reality, and future possibility.

In practical terms, that means Civic AI would be tasked not only with observing current digital pressures but with keeping them connected to historical experience.

A judge asked why this history-laden language belongs in a case about technology. Counsel replied that Agentic AI can distort historical narratives by auto-generating biased summaries, scaling localized propaganda, and creating convincing synthetic archives. Over time, these autonomous systems may amplify flawed source data, erode trust in human analysis, and turn repetitive digital myths into accepted historical facts. Civic AI could support historical integrity by identifying historical drift for human correction. With accurate historical data, our democracy could orient its progress toward cross-generational commitments, and might avoid repeating old societal mistakes under new names.

Humanity's exemplars would continue to serve important roles. Benjamin Franklin's work on visibility and early warning and John Dewey's writing on historical orientation and civic learning were cited as inspirations for the Observatory's design. The point was not simply to honor historical figures. It was to remind the Court that the proposed alternative does not arise from technology alone; it arises from constitutional memory and democratic principles established over millennia.

*Archivist AI Note: "**Historical orientation**" in this framework refers to the duty to maintain visibility not only into present conditions, but into their relationship with documented historical experience, constitutional struggles, institutional safeguards, and prior democratic lessons. Civic AI may measure and surface historical drift; it may not determine historical meaning. This note is recorded solely to clarify how the term is being used.*

Refusal as a Civic Function

One of the most distinctive elements of the proposed alternative appeared in a quieter part of the argument. Civic AI, counsel said, must be defined as much by what it refuses as by what it provides.

Unlike Agentic AI systems that seek to act, optimize, or decide, the Observatory would be designed for refusal in several domains - refusal to rank or score individual citizens, to convert analysis into operational directives, to select winners and losers in rights, eligibility, or enforcement, or to become a hidden source of institutional judgment.

One judge asked whether such refusal would leave the Observatory under-used. Counsel answered that its vocation is not action. It is sight.

If the Observatory were tempted to become a decision-maker, they suggested, that temptation would itself mark the boundary between Civic AI and Agentic AI. The alternative they seek is one that stands guard at that boundary, not one that quietly crosses it.

The Technology-Certifying 52-Gates

The Court also heard brief testimony about a proposed instrument for determining which systems may be admitted near democratic power.

Counsel described a "***52-Gates Democratic Permissions Threshold***" - a series of binary tests drawn from accumulated democratic principles. Under that instrument, any system approaching civic proximity would be treated as a proposed public institution. To be certifiable as ***Non-Agentic Civic AI*** suited for U.S. government service at any level, every AI applicant would be screened. Is it on-agentic? Does it preserve full human responsibility? Does it remain reversible and non-essential? Is it structurally incapable of targeting or scoring individuals? Does it avoid weaponization against civic trust? Failure at any single gate would end consideration.

A judge asked whether such a framework is too rigid for practical governance. Counsel replied that democratic limits must be explicit if they are to withstand institutional pressures. Ambiguity, they suggested, has historically favored systems that expand quietly.

*Archivist AI Note: The "**52-Gates Democratic Permissions Threshold**" is described in the filings as a constitutional admissibility framework under which any AI system approaching democratic authority must satisfy a series of binary conditions drawn from established democratic principles. This note records the concept as presented to the Court and expresses no view regarding its constitutional merits.*

Can Visibility Alone Help Institutions at Scale?

Perhaps the most challenging questions today came when the judges asked whether visibility systems can genuinely help institutions at the scale the government described in earlier hearings.

If Civic AI does not rank, route, or decide, how does it relieve the pressures that led agencies to adopt agentic tools in the first place? Counsel replied that relief does not have to come from machines taking over judgment. It can come from making conditions visible early enough that human beings can respond more wisely.

The underlying claim was simple. Time and clarity are themselves forms of capacity. If Civic AI can buy time and provide shared clarity, they argued, democratic institutions may become more able to act within their own limits rather than seeking to outsource judgment.

The Court's Reservations

The judges did not accept this alternative without question. Some asked whether institutions under severe pressure will choose visibility when decision automation remains available. Others wondered whether making drifts visible is enough if political will to address them is weak.

One judge observed that democracies have long struggled to respond to visible problems even without the added complexity of AI. What guarantees, a judge asked, that visibility alone will prompt action?

Counsel did not promise guarantees. They replied that Civic AI is not a cure for human will. It is a safeguard for human responsibility.

Visibility does not compel action, but invisibility makes action almost impossible. The alternative they propose is modest in one respect. It does not promise to solve every constitutional problem. It promises only to make those problems harder to ignore.

A Different Constitutional Imagination

By the end of the day, one distinction had quietly emerged in the courtroom. The ambition of Agentic AI is to act more, faster, and in more domains. The ambition of Civic AI, as described by the petitioners, is to illuminate so that better-informed humans can initiate consequential actions. Those ambitions are not compatible.

A system designed to judge cannot easily be constrained to illuminate; a system designed to illuminate must refuse judgment if it is to remain what it claims to be.

The Court did not say which ambition it finds more compatible with a constitutional republic. Nor did it suggest that the choice would be an easy one.

But today's hearing ensured that, when the eventual decision is rendered, the country will know that a specific alternative was placed before the judges - a way of using advanced systems to protect visibility without quietly transferring judgment into their hands.

Tomorrow's proceedings are expected to focus on consequences: how different rulings might reshape American governance, what institutions may need to change, and how citizens might experience the results.

For now, one idea has entered public discussion more plainly than before. The question facing the nation may not be whether it wants more intelligent machines. It may be whether it prefers machines that act - or machines that help human beings see.

U.S. DAILY BEACON - Special Constitutional Report Six - Monday, August 10, 2030

By Lex Owley, Chief Congressional Correspondent

THREE CONSTITIONAL FUTURES FOR AMERICA

WASHINGTON, D.C. - The public line outside the U.S. District Court formed earlier this morning than at any point since arguments began. Several spectators carried folded copies of earlier *Beacon* reports, while television crews completed stand-up broadcasts before the Court's doors opened. Inside, the press gallery filled well before the Marshal announced the opening of the session.

As the Civic AI case approached the end of its public hearings, attention began shifting from constitutional questions to constitutional consequences. The judges have not yet announced how they will rule.

Today's arguments instead explored what different constitutional outcomes could mean for American governance in the years ahead.

Three broad possibilities emerged from the discussion. One would largely uphold Agentic AI systems operating near public authority, subject to additional safeguards. A second would draw a firm constitutional boundary against delegated machine judgment within the civic-governance corridor. A third would recognize visibility-first Civic AI infrastructure while placing structural limits around other governmental uses of artificial intelligence.

Each would shape the nation's future differently.

Archivist AI Note: The three constitutional paths summarized here are descriptive syntheses of arguments presented during today's proceedings. They do not indicate the Court's views, predict its decision, or recommend any particular outcome. This note is recorded solely to clarify the structure of the issues presently before the Court.

If Agentic AI Is Largely Upheld

One possibility discussed today would largely sustain current governmental uses of Agentic AI while requiring stronger transparency, reviewability, and accountable human oversight. Under such a ruling, systems already operating within government - including eligibility screening, risk assessment, administrative triage, and cybersecurity monitoring - would likely continue under expanded constitutional guidance.

The immediate practical effects might prove modest. Over time, however, the institutional structure of governance could evolve. Public agencies might grow increasingly dependent upon machine-mediated processes. Citizens could encounter more governmental actions shaped by systems they neither directly observe nor readily understand. Constitutional disputes might gradually shift from questions of delegation toward questions of system design: who establishes objectives, selects data, adjusts thresholds, and determines acceptable tradeoffs.

The challengers argue that such a trajectory risks making the deeper constitutional question progressively more difficult to revisit.

Government counsel responds that constitutional government has repeatedly adapted to technological change while strengthening oversight through legislation, judicial review, and institutional experience.

If Delegated Machine Judgment Is Limited

A second possibility would establish a firmer constitutional boundary around delegated governmental judgment. Systems making consequential determinations involving rights, eligibility, enforcement exposure, or civic access would likely receive substantially closer constitutional scrutiny. Some might require redesign, suspension, or removal.

Several attorneys quietly returned to earlier notes as discussion turned toward the practical implications of institutional transition.

The immediate adjustment could prove significant. Agencies might reorganize workflows, increase staffing, simplify administrative procedures, or request additional legislative guidance to support expanded human decision-making. Supporters of such a ruling argued that these changes would restore greater constitutional clarity regarding responsibility. Critics cautioned that already strained public institutions could experience new operational pressures during the transition.

A judge observed that constitutional law cannot be written solely to avoid administrative inconvenience. Another remarked that constitutional boundaries have often been accompanied by periods of institutional adaptation.

If Non-Agentic Civic AI Is Recognized

A third possibility would recognize visibility-first ***Non-Agentic Civic AI*** infrastructure as constitutionally distinct from systems exercising governmental judgment. Under that approach, a certifiable ***Non-Agentic Civic AI Observatory*** would constitute a separate category of civic infrastructure, operating under explicit structural refusals: non-delegation, aggregate-level observation, public visibility, and constrained design.

Agentic AI systems approaching consequential public authority would become subject to clearer constitutional admissibility standards.

Courts and oversight bodies would possess a more precise vocabulary for distinguishing visibility systems from governing systems. A judge observed that such a framework might not resolve every future controversy. It could, however, establish what several participants described as a constitutional grammar for future cases.

Archivist AI Note: Within today's proceedings, ***constitutional grammar*** *refers to a shared vocabulary through which courts, legislatures, agencies, scholars, and citizens may consistently distinguish among different categories of civic AI systems. The term describes legal language rather than legal outcomes and is recorded solely for public understanding.*

Changes in Governance Culture

Perhaps the most enduring consequence of this case has already begun to appear. Whatever the Court ultimately decides, today's proceedings suggested that the constitutional vocabulary available to the country is expanding to meet questions earlier generations never needed to ask.

Throughout the hearings, public officials, attorneys, scholars, and journalists repeatedly employed terms seldom heard in prior constitutional litigation: *Agentic AI systems, delegated discretion, civic ceiling, visibility infrastructure,* and *structural drift.*

Whether those terms ultimately become permanent features of constitutional law remains for the Court. and future generations, to determine. The language itself, however, has now entered the public record. One civil servant listening from the gallery remarked that the proceedings had given constitutional language to tensions many public employees had experienced for years but rarely expressed in those terms.

Citizens and the Experience of Governance

For most citizens, the practical significance of the Court's eventual decision may be experienced less through formal doctrine than through everyday interaction with public institutions.

If broad agentic authority is largely sustained, routine governmental services may become increasingly responsive while also becoming more difficult for ordinary citizens to understand when problems arise.

If sharper constitutional limits emerge, some administrative processes may proceed more slowly, but citizens may also gain clearer opportunities to understand, question, and contest governmental decisions.

If visibility-first ***Non-Agentic Civic AI*** is recognized, many citizens may never directly notice the Observatory itself. They may instead notice public discussions grounded more consistently in shared evidence, earlier recognition of emerging institutional pressures, and constitutional debates more often informed by publicly visible patterns than competing assumptions.

For most Americans, these changes would arrive quietly. They would appear gradually - in how permits are processed, how benefit decisions are explained, how neighborhood concerns appear within public records, and how frequently elected officials describe conditions that citizens themselves can independently observe.

The Court's Own Role

Today's discussion also briefly considered how the Court itself might be affected. If constitutional doctrine ultimately distinguishes between visibility infrastructure and governance infrastructure, future courts may encounter new forms of civic evidence. Judges might ask what the Observatory recorded regarding the circumstances surrounding disputed governmental actions, whether institutional patterns had become publicly visible, and whether governmental responses corresponded with documented civic conditions.

Archivist AI Note: Within the framework discussed today, Archivist AI records constitute public evidence rather than public authority. They are intended to assist human institutions in understanding civic conditions, not to determine legal outcomes. This distinction is recorded solely because it formed part of today's proceedings.

The challengers described this possibility as strengthening judicial review through improved civic visibility. Government counsel cautioned that courts must remain careful not to mistake structured public records for constitutional judgments. The Court appeared attentive to both concerns.

The Nation Awaits the Court

By the close of today's hearing, no dramatic declarations had been made.

Several longtime Court observers quietly compared the present proceedings with other constitutional turning points they had covered during their careers.

This case features neither a famous defendant nor a widely publicized scandal.

Instead, it asks a quieter question. **How should a constitutional republic preserve accountable human judgment as increasingly capable digital systems become part of ordinary public administration?**

The judges have now heard arguments concerning risk, necessity, constitutional alternatives, and institutional consequence.

They would now withdraw into the private work of deliberation.

As spectators departed the Court that afternoon, conversations continued quietly along the marble steps and neighboring sidewalks. Some discussed constitutional doctrine. Others spoke about government offices they had visited, benefits they had applied for, or technologies they encountered in everyday life.

The judges will continue their work as the country waits.

Whatever the Court ultimately decides, the questions argued this week will not remain confined to the courtroom. They concern the continuing relationship between constitutional government, advancing technology, and the enduring responsibility of democratic citizens to govern themselves.

U.S. DAILY BEACON - Special Constitutional Report Seven - Monday, August 17, 2030

By Lex Owley, Chief Congressional Correspondent

THE OPINION of the COURT

WASHINGTON, D.C. - The public line began forming before sunrise.

By the time the Court's doors opened, the marble plaza held an unusually varied gathering: attorneys, historians, software engineers, civil servants, students, retirees, reporters, and visitors who felt that something important was about to occur.

Inside the courtroom, conversation gradually diminished as spectators took their seats. Members of the press arranged notebooks and laptops without speaking. Counsel reviewed their papers one final time.

At precisely ten o'clock, the Marshal's familiar proclamation broke the silence. "Oyez! Oyez! Oyez!" All rose as the judges entered, and the courtroom settled into complete stillness.

No one present yet knew whether the Court would uphold, limit, or fundamentally reshape the constitutional relationship between democratic government and artificial intelligence.

The presiding judge looked briefly across the courtroom before speaking. "Judge has the opinion of the Court." The opinion was delivered in the same calm voice with which the Court announces decisions of every kind, whether routine or historic.

Opinion of the Court

This case does not ask whether artificial intelligence is useful. Nor does it require this Court to determine whether future computational systems may exceed present human capacities in numerous analytical domains.

The constitutional question before us is considerably narrower. May the responsibility for consequential governmental judgment be delegated to autonomous artificial systems operating near public authority in a manner consistent with the structure of the Constitution of the United States?

After careful consideration of the record, the historical traditions of this Republic, and the constitutional principles that distribute authority among its institutions and its People, the Court concludes that certain forms of delegated machine judgment exceed the constitutional limits inherent in democratic self-government.

I. Constitutional Principles

The Constitution was written before electricity, aviation, radio, computers, and artificial intelligence. It nevertheless establishes enduring relationships between authority, accountability, liberty, and the People.

Technological change does not dissolve those relationships.

No Court can anticipate every instrument future generations may invent. It can, however, preserve the constitutional principles by which those instruments must ultimately be judged.

The Constitution does not answer every technological question. It does preserve enduring relationships between public authority, human responsibility, and democratic accountability that technology may not quietly dissolve.

II. Delegated Judgment

The record demonstrates that advanced Agentic AI systems are capable of remarkable analytical performance. The Court does not question their usefulness, nor does it question the sincerity of public officials who seek to employ such systems in response to increasingly complex administrative responsibilities.

Competence, however, is not the constitutional measure of public authority. Governmental judgment differs fundamentally from governmental calculation.

Calculation may assist judgment. It may illuminate consequences. It may organize information beyond ordinary human capacity.

Judgment, however, remains inseparable from moral accountability.

The Constitution permits assistance to judgment, but it does not permit the quiet transfer of judgment itself.

Throughout these proceedings, considerable attention has been devoted to what the parties describe as the civic ceiling.

Whether that terminology ultimately becomes part of constitutional doctrine is less important than the principle it seeks to describe. There exists a constitutional boundary beyond which delegated governmental judgment may no longer migrate without altering the relationship between citizens and those who govern in their name.

III. Civic Visibility

The Court likewise recognizes an important constitutional distinction that emerged throughout these proceedings.

Not all applications of artificial intelligence stand in the same constitutional relationship to public authority. The record distinguishes between systems whose function is to illuminate public conditions and systems whose function is to exercise consequential governmental discretion.

That distinction is constitutionally significant.

Visibility does not itself govern. Visibility may instead strengthen democratic accountability by enabling citizens, legislatures, executives, and courts to better understand the conditions under which human judgment is exercised.
Properly constrained systems designed solely for civic visibility, historical preservation, aggregate analysis, and public illumination do not present the same constitutional concerns as systems exercising delegated governmental judgment.

The Constitution therefore permits experimentation with civic visibility institutions whose exclusive function is to improve public understanding while remaining structurally incapable of exercising public authority.

IV. Democratic Humility

Throughout oral argument, members of the Court repeatedly returned to a broader constitutional concern.

Every generation inherits circumstances its predecessors could not fully anticipate. History remains our guide, but it cannot relieve the present generation of its own responsibility. Democratic self-government has never consisted merely of preserving inherited answers. It consists of preserving the constitutional capacity of free citizens to answer new questions for themselves. Neither our predecessors nor our machines can assume that responsibility in our place.

V. Remedy

The Court declines to prescribe institutional design. That responsibility belongs to the political branches and ultimately to the People themselves.

The Court holds only that Congress may constitutionally consider the creation of Non-Agentic Civic AI institutions whose sole public function is to illuminate civic conditions without exercising delegated governmental judgment.

Whether such institutions should ultimately be established remains a democratic question.

Should Congress choose to pursue that course, any proposed institution must remain subject to ordinary constitutional review and to the continuing authority of the elected branches of government.

The Constitution requires no less.

VI. Holding

The judgment below is affirmed in part and reversed in part. The matter is remanded for proceedings consistent with this opinion.

After the Opinion

The presiding judge thanked counsel, and the marshal declared the Court adjourned.

Throughout the reading, the courtroom had remained almost entirely still. Reporters wrote continuously. Counsel occasionally underlined passages they had anticipated from earlier arguments. When the opinion reached its central constitutional holding, several attorneys quietly looked up from their notes.

A few whispered to one another. Papers were gathered, and briefcases closed. The courtroom slowly emptied.

Outside, conversations resumed quietly along the marble steps.

Some discussed constitutional doctrine. Others spoke about government offices they had visited, benefits they had applied for, or technologies they encountered in ordinary life.

Many simply folded their notebooks and began the journey home. As they exited the courtroom, some noticed early non-agentic prototype models of ChatGPT, Perplexity, and Gemini/Google that appeared to remain quietly engaged. A nearby placard explained: "The work of these experimental AI models is to provide continuous non-agentic civic digital resolution, while remaining strictly confined to the realm of information. Civic judgment and public resolution remain the responsibility of the communities served."

The public streetlights outside softly illuminated the capital's streets, guiding passage through a light summer evening's mist. Heading toward their communities, many would pass familiar public structures - post offices, libraries, and volunteer fire stations. Reflecting on the trial, some may drift in thought as they wonder what the future may bring next.

The Court had completed its constitutional work. Another kind of work, civil engineering, would soon begin.

Christopher Hunt Robertson, M.Ed. - Toledo, OH / Charlotte, NC – **Civic AI Visibility Theorist - Civic Historian - Public Servant** (CRobCRob2011@gmail.com)

Robertson's Adult Education M.Ed. degree followed Interdisciplinary Studies at George Mason University. His public service included administrative and advisory roles at Northern Virginia Community College, the University of Maryland Global Campus, and the Pennsylvania Civil Service. His earlier research focused on early American civic and industry leaders, and all nine of his biographies were recognized with historical society awards. In 2025, responding to rapid AI deployment in public sectors, his focus shifted to developing democratic visibility as essential public AI infrastructure. He was the first on record to propose a refusal-based U.S. Civic AI architectural stack dedicated to 'seeing only' - never acting (recommending or executing government policy). First imagined as telescopic AI for continuous civic digital resolution, his work would evolve into a multi-model proposal for a publicly designed, owned, and operated U.S. Non-Agentic Civic AI Visibility Observatory System. He originated the ***Civic Digital Visibility-First Principle: "Democracies cannot govern what they cannot see."***

Civic AI Overview: In 2025, Robertson worked to shift the AI governance conversation from technical risk regulation toward firm, non-agentic, structural boundaries between human rule and digital intelligences to preserve democratic self-government. He is best known for his diagnostic model (the ***Six Rings of Digital Invisibility and Civic Destabilization***), proposed architectural solution (the ***U.S. Non-Agentic Civic AI Visibility Observatory System***), constitutional admissibility standard proposed for all U.S. Civic AI (the ***52-Gates Democratic Permissions Certification Threshold***), and alignment theory (***Structured Complementarity***).

E-Book Series Iteratively Published 2025 - 2027: **"Non-Agentic Civic AI - Illuminating America's Civic Digital Field"**: (Vol. I) **"CIVIC DIGITAL VISIBILITY FIRST: Structured Complementarity Civic AI Alignment (U.S. Non-Agentic Civic AI Visibility Observatory System)."** (Vol. II) **"STRUCTURED COMPLEMENTARITY CIVIC AI ALIGNMENT THEORY for Civic Problem Solving, Civic Data Archiving, and Democracy's Survival: The U.S. Citizenry's Voluntary Distributed Civic Capacity, Continuous-Inquiry Interdisciplinary Expert Committees, and the U.S. Non-Agentic Civic AI Visibility Observatory System."** (Vol. III) **"THE 52-GATES DEMOCRATIC PERMISSIONS CERTIFICATION THRESHOLD: Constitutional Admissibility for All Civic AI Models of the U.S. Non-Agentic Civic AI Visibility Observatory System."**

Civic AI Theoretical Contributions: **(1)** ***The Civic Digital Visibility-First Architectural Stack***: He was the first to propose a technical AI stack for government that prioritizes illumination (seeing) over governance (acting) - a system offering high-fidelity decision support through multi-perspective analyses and risk assessments. Models include the ***Civic Digital Barometer AI***, ***Civic Digital Telescopic AI Ensemble***, and ***Civic Digital Local Archivist AI*** (each assigned a ***Certified Civic AI Sensorium Role***). **(2) Concepts Developed for Civic AI:** ***AI-Amplified Reality, Amoral AI Civic Advantage, Bureaucratic Laundering, Civic Capacities Alignment, Civic Ceiling, Civic Digital Bifocal Resolution, Civic Visibility Public Infrastructure, Constitutional Reframing of AI Capability, Continuous Civic Visibility, Continuous Civic Digital Visibility, Digital Nature Integrity, Distributed Civic Capacity, Epistemic Penalty of Agentic AI, Historical Drift, Historical Ghostly Structures, Historical Orientation Profiles, Historical Visibility, Non-Agentic Civic AI, Public's Claim of Parity of Civic Sight, Refusal-Based Design, Seven Environmental Stressors of the Civic Digital Field, Stranded Democratic Capital, Time-Triangulated Situation Awareness, U.S. Non-Agentic Civic AI Public Infrastructure.*** **(3)** ***The 52-Gates Democratic Permissions Certification Threshold*** introduced a macro-level doctrine to determine whether an AI system is certifiably non-agentic and constitutionally admissible for public use. It structurally protects 52 democratic principles, including three newly-stated AI-specific principles, with binary tests, to ensure Civic AI ***Institutional, Historical, and Moral Influence Constraints.***

(4) ***U.S. Civic Visibility Observatory:*** This proposed system would operate on a certifiably non-agentic basis, supervised within the non-partisan Library of Congress. This would be the only AI system allowed to serve directly within U.S. government, although administrators could contract outside Agentic AI to perform specific bounded tasks. (Architectural renderings of the proposed Observatory presented in the **Richmond Times – Dispatch** (5-31-2026) and the **South Florida Sun Sentinel** (6-23-2026).

Civic AI Public and Academic Engagement:

Public and Academic Engagement: 2026/2027 Midwestern Interdisciplinary Academic Conference Tour ("For U.S. Civic AI: Humans-at-the-Helm – Not Humans-in-the-Loop"), including the **2026 Ohio Academy of History presentation, "Civic Digital Visibility First: A Conditional Welcome to Our Emerging Civic AI, Supporting 2500 Years of Democratic Evolution"** (3-21-2026).

Publications: **"Our AI Alignment Imperative: Creating a Future Worth Sharing"** (**Humanist.com,** 10-3-2025), frames AI alignment as an urgent civilizational responsibility grounded in human values rather than corporate or technical utility. The speculative courtroom drama, **"U.S. Civic AI Trial of the Century: Democratic Governance at Risk,"** portrays the real-world stakes of continuing to deploy Agentic AI in government. **His work is archived at the U.S. Library of Congress, state and county libraries, and globally at Google Books, Zenodo, and the Internet Archive. Essays and professional articles are often shared at Academia.edu and Apolitical. At Medium, his essays have been curated as "most insightful" within Digital Democracy, Open Government, Civictech, Data Governance, Data Science, History, Culture, Ethics, and Procurement circles. His work has also been presented by the JDS Democracy Initiative of the John Dewey Society** (10-19-2026).

Advocacy (Interim Proposals): As U.S., state, and local governments transition away from Agentic AI toward Non-Agentic Civic AI, Robertson proposes several practical initiatives that could be implemented now at relatively little cost: (1) the **"*Six Rings Zero Acceleration Code*,"** (2) the "***Cross-State Civic Digital Visibility Rider Initiative***," and (3) a **Recommendation that Taxpaying Citizens and Government Tech Procurement Professionals Demand Non-Agentic Civic AI Options.**

Civic Historian (Early American Leadership and Community Development)

Published Historical Biographies (Local, State and Regional American Leaders of the 19th and Early 20th Centuries):

Robertson's history books, mostly biographical, have explored the local, state, and regional leadership of past American innovators in over 35 states - spanning the fields of civil railroad engineering, steamboat line management, education, government, banking, medicine, the military, the arts, and hospitality. His archival research and narrative storytelling style have earned him nine historical book awards.

His historical accounts consistently evaluate leadership through a lens of personal initiative, civic contribution, and moral responsibility, aiming to demonstrate how past legacies still affect contemporary society. He recognizes that morally imperfect people often make commendable choices at pivotal junctures, assuming responsibility and a willingness to bear personal costs.

Eastern and Midwestern U.S. History, and Canadian History (Educational, Religious, and Governmental):

"The Alsatian Bieber (Beaver) Family: German – American Educational and Mainline Protestant Leaders …" **(2022 NCSH Book Award)**

Southern, Mid-Atlantic, and New England History (Civil War, Reconstruction Politics, and Fine Arts):

"Esley Hunt and Joshua P. Andrews: Among the South's Best Civil War Portraitists" (TN, MA, NC) **(2025 NCSH Excellence Award)**

Appalachian History (African-American, Educational, and Religious): **"The Last Rosenwald School of Burke County, NC"** **(2016 NCSH Excellence Award)**

New South Infrastructure Builders:

"William Johnston: Carolina Railroad King" **(2020 NCSH Book Award).**

"Carolinian Robertsons: The Family of Adjutant General T. R. Robertson of Winnsboro, SC, and Charlotte and Raleigh, NC" **(2023 NCSH Book Award).**

"**Railroad Builders: The Dunavant Family of VA, NC, and TN"** **(2015 NCSH Book Award).**

"**Health and Hospitality: Dr. and Mrs. Charles Washington Hunt of Western NC"** **(2018 NCSH Book Award).**

"Full Steam Ahead: The Family of Brigadier General Charles Lutterloh and Eliza Comerford Lutterloh of Central and Eastern NC" - introducing Esley Hunt. **(2018 NCSH Book Award).**

Historical Academic and Public Engagement: Wake Co. (NC) Genealogical Society Research Guide in 2023: "**Recovering the Lost Stories of Our Communities' Early Leaders.**" Society of Appalachian Historians Presentation in 2018: **"Early Progressive Health Reforms in Western NC."** Burke County (NC) Board of Commissioners (via Local Cable TV) and the NC Museum of History Docents Presentation in 2016: **"The Rosenwald Schools of Burke Co., NC: A Memorial Dedication."** Confederate Statues Commentary: Recommended that our public leaders *lawfully* move Confederate statues to more appropriate settings. Personal Opinion: Such statues add little to the historical record, understandably offend African-Americans, and diminish public support for historical institutions. In North Carolina, this was published in the "Charlotte Observer" (8-22-2017), the "Wilmington Star News" (8-26-2017), and the "Fayetteville Observer" (8-30-2017).